

Cloud Based Collaborative Road Surface Monitoring using Deep Learning and Smartphones

Final Report

by

Yunyi Jia

yunyij@clemson.edu

Clemson University

Gurcan Comert

Gurcan.Comert@Benedict.edu

Benedict College

Akshatha Ramesh, Dhananjay Nikam, Venkat Narayanan Balachandran,
Longxiang Guo, Rongyao Wang, and Leo Hu
Clemson University

July 2023



Center for Connected Multimodal Mobility (C²M²)



Benedict College



THE
CITADEL
THE MILITARY COLLEGE OF SOUTH CAROLINA

SCState
UNIVERSITY



UNIVERSITY OF
SOUTH CAROLINA

200 Lowry Hall, Clemson University
Clemson, SC 29634

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, by the Center for Connected Multimodal Mobility (C²M²) (Tier 1 University Transportation Center) Grant, which is headquartered at Clemson University, Clemson, South Carolina, USA, from the U.S. Department of Transportation's University Transportation Centers Program. However, the U.S. Government assumes no liability for the contents or use thereof.

Non-exclusive rights are retained by the U.S. DOT.

ACKNOWLEDGMENT

This study is partially supported by the Center for Connected Multimodal Mobility (C2M2) (USDOT Tier 1 University Transportation Center) headquartered at Clemson University, Clemson, SC. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of C2M2, and the official policy or position of the USDOT/OST-R, or any State or other entity, and the U.S. Government assumes no liability for the contents or use thereof.

Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle Cloud-based Collaborative Road Condition Monitoring using In-Vehicle Smartphone Data		5. Report Date
		6. Performing Organization Code
7. Author(s) Yunyi Jia; ORCID: 0000-0003-1334-1384 Gurcan Comert; ORCID: 0000-0002-2373- Akshatha Ramesh; ORCID: 0009-0004-5261-0211 Dhananjay Nikam; ORCID: 0009-0008-0774-3268 Venkat Narayanan Balachandran; ORCID: 0000-0001-8798-8493 Longxiang Guo; ORCID: 0000-0002-6889-6099 Rongyao Wang; ORCID: 0009-0004-0314-6045 Leo Hu; ORCID: 0000-0003-2657-5382		8. Performing Organization Report No.
9. Performing Organization Name and Address Clemson University, 4 Research Drive, Greenville, SC 29607 Benedict College, 1600 Harden St, Columbia, SC 29204		10. Work Unit No.
		11. Contract or Grant No. 69A3551747117
12. Sponsoring Agency Name and Address Center for Connected Multimodal Mobility (C ² M ²) Clemson University 200 Lowry Hall, Clemson, SC 29634		13. Type of Report and Period Covered Final Report July 2023
		14. Sponsoring Agency Code
15. Supplementary Notes		
16. Abstract <p>Ensuring the safety of transportation systems requires monitoring the conditions of roads. Traditional monitoring and inspection of road conditions require surveyors to walk or drive along the roads to search for defects manually. Such processes require a lot of human and equipment efforts, which however can still hardly provide real-time information on road conditions. Existing automated road condition monitoring approaches usually require special vehicles equipped with specific sensors and corresponding processing and computing devices. In addition, these approaches only use one single vehicle to perform the detection on its own and the vehicle usually still needs to be driven by a surveyor. Therefore, in this project, we developed a much more cost-effective approach to monitoring road conditions by cloud-based collaborative monitoring using in-vehicle smartphones which could come from any public vehicle user.</p> <p>When a vehicle drives over a certain type of road defect, the acceleration signal, especially the vertical acceleration, will have a unique pattern in the trajectory. The type of road defect can be identified from the general shape of the acceleration wave. Meanwhile, the amplitude of the wave reflects the vehicle speed and the severity of the defect. In this project, we trained a Long Short-Term Memory (LSTM) based deep learning network to complete the identification of defect types using acceleration data. Sometimes LSTM could have difficulty in deciding the defect type solely based on accelerations since the smartphone would be placed in the passenger cabin, and the motion it measures will be filtered by the vehicle suspension. Thus, we trained YOLO (You Only Look Once) deep learning network to detect and identify defects from the live video taken by the smartphone's camera. We then fused the road condition detection results of multiple deep learning approaches from smartphones of multiple vehicles in order to get holistic monitoring of the road condition. The data including the smartphone motion and vision-based road condition detection results and the GPS locations of the vehicles would be sent to a cloud server through cellular networks. All detection results were then fused with the k-means clustering method based on their GPS locations, and the top three most occurred types of damage within a cluster were found to represent the road condition of that location.</p> <p>We developed a data collection app to collect acceleration and vision data from smartphones mounted on the windshield of multiple cars. The data used for this experiment was collected over various roads of Greenville, Spartanburg, Clemson, and Columbia area in South Carolina, USA. Eventually, we were able to get an accuracy of 94% from the trained LSTM model, and 87.5% from the trained YOLO in classifying potholes, cracks, and normal road surfaces. We also created a web page that displays the fusion results of detected road damage on a map. The web page enables concerned authorities to view the road damages reported by the users with the help of our developed mobile application.</p>		

17. Keywords		18. Distribution Statement	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages	22. Price NA

Table of Contents

DISCLAIMER	ii
ACKNOWLEDGMENT	iii
EXECUTIVE SUMMARY	1
CHAPTER 1 Introduction	2
CHAPTER 2 Literature Review	3
2.1 Related Work	3
2.2 Challenges and Gaps	4
CHAPTER 3 Cloud-Based Collaborative Road-Surface Monitoring	5
3.1 Framework of Cloud-Based Collaborative Road-Damage Monitoring Method...	5
3.2 Motion-Based Road-Damage Detection	6
3.3 Vision-Based Road-Damage Detection	6
3.4 Cloud-Based Collaborative Fusion	7
3.5 Road-Damage Severity Estimation	9
CHAPTER 4 Results	12
4.1 Experimental Setup and Data Collection	12
4.2 Experimental Result and Analysis	13
CHAPTER 5 Conclusions	22
REFERENCES	23

List of Tables

Table 1: Count of data across road-damage types.....	13
Table 2: Count of acceleration data across classes.....	13
Table 3: Count of images across classes after augmentation	13
Table 4: Accuracy results of training data sets	15
Table 5: Accuracy results of testing data sets.....	15
Table 6: Tuning parameters of LSTM	17
Table 7: Road-damage-severity estimation error.....	21

List of Figures

Figure 1 Cloud-based collaborative road damage monitoring using in-vehicle smartphone data and deep learning.	5
Figure 2 LSTM model structure for motion-based road damage detection.....	6
Figure 3 Architecture of cloud-based collaborative fusion.	7
Figure 4 Cloud-based collaborative fusion flowchart.	8
Figure 5 Back-end process of cloud-based collaborative fusion.....	9
Figure 6 Data visualization of different road-damage estimation parameters.....	10
Figure 7 Road-damage severity estimation based on acceleration data in the worst-case scenario.	10
Figure 8 Cloud-based road-damage severity estimation process.....	11
Figure 9 Screen view of the data-collection app	12
Figure 10 Data-collection example.....	12
Figure 11 YOLOv5 mean average precision over epochs	14
Figure 12 YOLOv5 loss over epochs on the training and validation data set.	14
Figure 13 Confusion matrix—YOLOv5 model on the training data set.	15
Figure 14 Confusion matrix—YOLOv5 model on the testing data set.	16
Figure 15 YOLOv5 precision-recall curve.	16
Figure 16 YOLOv5 inference on validation images.....	17
Figure 17 YOLOv5 inference on validation images.....	18
Figure 18 LSTM classification loss over epochs.....	18
Figure 19 Confusion matrix—LSTM model on the training data set.	19
Figure 20 Confusion matrix—LSTM model on the testing data set.....	19
Figure 21 User interface web page.....	20
Figure 22 Damage-severity estimation validation.	21

EXECUTIVE SUMMARY

Ensuring the safety of transportation systems requires monitoring the conditions of roads. Traditional monitoring and inspection of road conditions require surveyors to walk or drive along the roads to search for defects manually. Such processes require a lot of human and equipment efforts, which however can still hardly provide real-time information on road conditions. Existing automated road condition monitoring approaches usually require special vehicles equipped with specific sensors and corresponding processing and computing devices. In addition, these approaches only use one single vehicle to perform the detection on its own and the vehicle usually still needs to be driven by a surveyor. Therefore, in this project, we developed a much more cost-effective approach to monitoring road conditions by cloud-based collaborative monitoring using in-vehicle smartphones which could come from any public vehicle user.

When a vehicle drives over a certain type of road defect, the acceleration signal, especially the vertical acceleration, will have a unique pattern in the trajectory. The type of road defect can be identified from the general shape of the acceleration wave. Meanwhile, the amplitude of the wave reflects the vehicle speed and the severity of the defect. In this project, we trained a Long Short-Term Memory (LSTM) based deep learning network to complete the identification of defect types using acceleration data. Sometimes LSTM could have difficulty in deciding the defect type solely based on accelerations since the smartphone would be placed in the passenger cabin, and the motion it measures will be filtered by the vehicle suspension. Thus, we trained YOLO (You Only Look Once) deep learning network to detect and identify defects from the live video taken by the smartphone's camera. We then fused the road condition detection results of multiple deep learning approaches from smartphones of multiple vehicles in order to get holistic monitoring of the road condition. The data including the smartphone motion and vision-based road condition detection results and the GPS locations of the vehicles would be sent to a cloud server through cellular networks. All detection results were then fused with the k-means clustering method based on their GPS locations, and the top three most occurred types of damage within a cluster were found to represent the road condition of that location.

We developed a data collection app to collect acceleration and vision data from smartphones mounted on the windshield of multiple cars. The data for this experiment was collected over various roads of Greenville, Spartanburg, Clemson, and Columbia in South Carolina, USA. Eventually, we were able to get an accuracy of 94% from the trained LSTM model and 87.5% from the trained YOLO in classifying potholes, cracks, and normal road surfaces. We also created a web page that displays the fusion results of detected road damage on a map. The web page enables concerned authorities to view the road damages reported by the users with the help of our developed mobile application.

CHAPTER 1

Introduction

One of the causes of U.S. road accidents today is road damage [1]. Road problems such as potholes and iced-over stretches of highways cause many traffic accident deaths each year. Well-maintained and monitored road surfaces could increase road-user safety, fuel economy, and comfort levels. The currently widely used methods for monitoring road-surface conditions include using profilometers [2] and surveying techniques. Surveying is the traditional road-condition monitoring and inspection method that requires a technician to walk or drive along the roads to search for defects manually. Such a process requires a lot of human and equipment effort and can still hardly provide timely information about road damages. In addition, the profilometer equipment is expensive and requires trained professionals to operate it, making this process costly and impossible to execute frequently.

To simplify monitoring processes by lowering the requirement of special equipment, smart detection and classification algorithms have been developed. A significant number of those methods are based on vehicle motion signals such as speed and acceleration. Another important category of those methods is vision-based detection. These recently developed methods are much more convenient than the traditional methods for detecting road damage, but they still rely on dedicated data-collection devices, which limits the scale of their applications. Thus, if an effective and cost-efficient method is developed, it would be beneficial to monitor road damage continuously to enhance the transportation system in terms of driving safety and comfort. The usage of smartphones in scientific research has increased exponentially in recent years. Smartphones come with increasing computing power and a great variety of sensors, such as accelerometers, gyroscopes, magnetometers, GPS, and cameras, and many people are carrying smartphones with them when they are driving. These features make using smartphones as monitoring nodes possible.

Thus, the objective and contribution of this project are to develop a cost-effective approach to monitoring road damage, including potholes and cracks, using in-vehicle smartphones. The proposed method consists of both vehicle-motion-based and vision-based road-damage detection and classification methods to improve detection accuracy. The method also contains a cloud-based fusion algorithm to fuse all road-damage detection results from different vehicles in order to provide holistic, accurate, and complete monitoring of road damage.

CHAPTER 2

Literature Review

Many research efforts have been devoted to the field of the smart monitoring of road conditions or road damage. This section will briefly review existing work that employs motion-based and vision-based detection methods, as well as cloud computing technologies.

2.1 Related Work

The mainstream of smart road damage/condition monitoring consists of motion-based methods that rely on vehicle speed, acceleration, and position data to make predictions on the status of the road surface. For example, [3] discussed a transfer function-based method to estimate power spectral density using the relationship between the road surface and vehicle acceleration and then estimate road-damage levels. The authors of [4] developed a piece of dedicated equipment called CarMote to monitor road-surface conditions. The device can measure accelerations and send data to roadside units. The authors of [5] designed a mobile sensor network using a public transportation system called BusNet. The system uses acceleration sensors mounted on buses to monitor the conditions of road surfaces. The authors of [6] used the statistics of vertical acceleration to detect potholes and bumps on roads. The authors of [7] applied thresholds to the value and the changing rate of vertical acceleration to detect potholes, and achieved a precision of up to 89%. The authors of [8] also used the changing rate of vertical acceleration to detect potholes, and the detection rate was around 80%. The authors of [9] trained a support vector machine (SVM) to detect traffic conditions based on a vehicle's longitudinal acceleration and obtained a false positive rate of 2.7%, and a false negative rate of 21.6%. The authors of [10] also employed SVM to detect potholes and bumps based on the standard deviation, mean, variance, and other features of a vehicle's vertical acceleration over 2 seconds; the false positive rate was 3%, and the false negative rate was 18%. In recent years, the usage of deep-learning models has become popular. The authors of [11] used both a convolutional neural network (CNN) and long-short term memory (LSTM) network to detect potholes, speed bumps, street gutters, and some other stability events using vehicle vertical acceleration and achieved an accuracy of 93% for CNN, and an accuracy of 82% for LSTM. The authors of [12] used both CNN and LSTM to detect the type of pavement based on the time and frequency domain features of vertical acceleration and achieved about 91% accuracy for LSTM, and 93% accuracy for CNN. This work also demonstrated that the placement of the inertial sensor (below suspension, above suspension, and near the dashboard) does not influence the detection accuracy by more than 2%. The authors of [11] analyzed which factors affect the acceleration measurements by smartphones in a moving vehicle the most and how they change road-roughness measurements.

Another major trend of existing road-surface condition monitoring is vision-based methods. The authors of [13] discussed using a CNN-based network for the detection of road-distress types using data from a specialized setup of cameras. The work by [14] evaluated the different segmentation algorithms for the pavement distress type. In [15], quantifying the distress type was considered along with identification and classification. Furthermore, the work by [16] discussed a deep fully convolutional model for crack detection (CrackPix), which leveraged well-known image-classification architectures for dense predictions by transforming their fully connected layers into convolutional filters. In recent years, the Road Damage Detection and Classification Challenge appeared and has drawn much attention. The authors of [17] used faster regions with convolutional neural networks (R-CNN) to detect eight different types of road damage in the challenge and achieved a mean F1-Score of 0.6225. The authors of [18] used a one-stage detector called "You Only Look Once" (YOLO-v4) for the challenge and achieved an F1 score of 0.628. The authors of [19] used a variant of Fast R-CNN, the bi-directional feature pyramid network (BiFPN), and

achieved an F1 score of 0.6455. Other participating teams also used variants of YOLO [20], [21] or R-CNN [22], [23], and their F1 scores were between 0.5 and 0.66.

Cloud computing technologies have also been used to help the monitoring of road damage or conditions. The authors of [24] offloaded heavy computation to a cloud database to achieve faster and more precise detection of road conditions on smartphones. The authors of [25] used road-condition data saved on the cloud to generate alerts for end users. The authors of [26] ran a fast unsupervised road-damage detection method on edge devices, and a slower but more precise machine-learning model on the cloud to realize fast, real-time detection and precise warning of road damage at the same time.

2.2 Challenges and Gaps

Traditional monitoring and inspection of road conditions require surveyors to go along the roads to search for defects. Such processes are very time-consuming, expensive, and labor-intensive. Existing automated road-condition monitoring research has been studying motion-based and vision-based detection approaches. However, such approaches usually require special vehicles equipped with specific sensors associated with complex processing systems, which is still time-consuming and expensive. In addition, most of these approaches cannot provide real-time and in-time monitoring of road conditions. Although some recent research efforts have preliminarily investigated the possibility of using smartphone sensors for a particular type of road-condition detection [11], they usually use just one type of sensors, such as a motion sensor or a vision sensor. Few have used both motion and vision sensors for detection. Furthermore, these approaches are all at the single-vehicle level, without the leverage of multiple vehicles for a more holistic and in-time detection. Furthermore, the identification of road-damage severity has not been investigated as a further detection of damage types. Therefore, a more time-efficient and cost-effective approach to holistically monitoring the real-time and in-time road conditions with damage severity is needed.

We propose a novel cloud-based collaborative road-damage monitoring approach using multi-sourcing in-vehicle smartphones in this project. It leverages a larger number of existing common vehicles on roads and uses both onboard motion and vision sensors from smartphones to conduct road-damage monitoring in a more cost-effective and time-efficient way. Furthermore, it leverages the multiple sources of detection from different vehicles with smartphones and fuses them in a cloud to achieve a more holistic detection of road conditions, including damage severity.

CHAPTER 3

Cloud-Based Collaborative Road-Surface Monitoring

This section introduces the details of the proposed cloud-based collaborative road surface monitoring method. The first subsection will present the framework of the proposed method. The second and third subsections will describe the vision-based road-condition detection using CNN, and the motion-based detection using LSTM. The fourth subsection will give details on the cloud-based fusion algorithm for detection results from different networks and vehicles. The fifth subsection will explain how the severity of the damage can be estimated using data on the cloud.

3.1 Framework of Cloud-Based Collaborative Road-Damage Monitoring Method

Modern cellphones are packed with advanced sensors, including a magnetometer, accelerometer, gyro, and GPS, which enable them to measure precise acceleration and speed. When vehicles are driven on roads with different surface conditions, cellphones will receive acceleration stimuli with different patterns and amplitudes, and such different patterns can be used by machine-learning models to detect and classify road damage such as potholes and cracks. Modern cellphones are also equipped with high-quality cameras that can directly 'see' the surface condition of a road when the vehicle is driven on it. Moreover, many people carry their cellphones every day when they travel, which makes mobile phones very suitable for forming a sensor network that monitors road damage continuously: as long as enough people are using the road-damage detection app that is going to be introduced in this report, the timeliness of road-damage information can be assured.

In this project, the proposed cloud-based collaborative road-damage monitoring method utilizes both motion and vision data from cell phones to achieve more reliable and precise detection of road-surface conditions. The framework of the method is shown in Figure 1. The method contains three major modules: the local detection and prediction module, the cloud-based fusion module, and the user interface (UI). The local detection and prediction module include dedicated deep-learning models for evaluating vehicle-motion input and vision input separately. The motion and vision-based detection results are shared to a cloud database through cellular communication. The cloud database will receive multiple detections from different vehicles passing the same area over time, and the cloud-based fusion module will fuse all those detections to generate a map that contains highly credible road-surface condition information. The UI will show users the road-damage map via a cellphone application and warn the driver of potential danger. Figure 1 is the overview of the proposed approach.

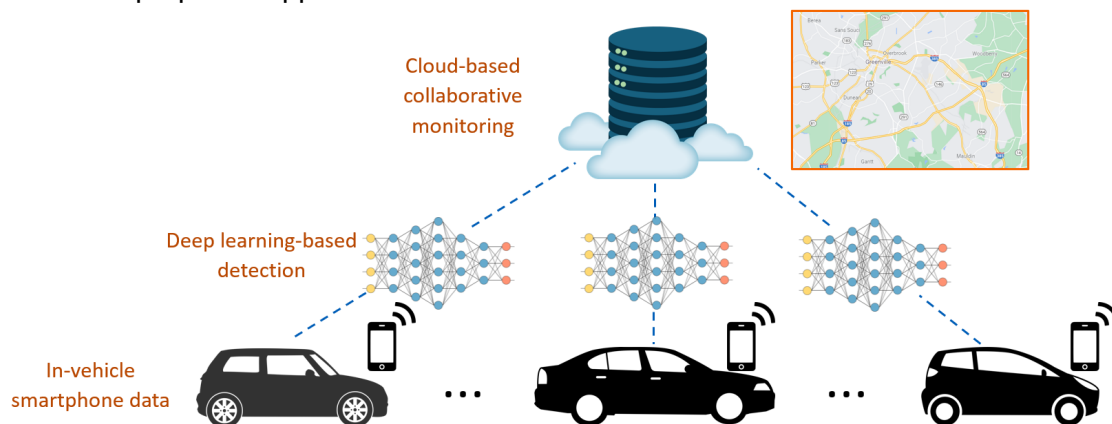


Figure 1 Cloud-based collaborative road damage monitoring using in-vehicle smartphone data and deep learning.

3.2 Motion-Based Road-Damage Detection

Sensors on smartphones such as accelerometers, gyroscopes, g-force, and magnetometers can reflect the profiles of the road when the smartphone is mounted in a car that is driven on the road. One of the recurrent neural network (RNN) variants, LSTM [27], is used in this project to utilize these motion data for road-surface classification. The reason for selecting the LSTM network is that, unlike standard feed-forward neural networks, LSTM has feedback connections. It can process single data points (such as images) and entire sequences of data (such as speech or video). This specialty is critical for our method, since a single acceleration or speed data point cannot indicate a certain type of road-damage condition. The data are meaningful only when they are in context.

The proposed method utilizes an architecture that is a stack of two fully connected layers and two LSTM layers. Figure 2 shows the architecture of the proposed architecture. The first fully connected layer has an input size of 200 hidden units, followed by another hidden layer. The two LSTM layers follow this, with 64 units each stacked on each other. An LSTM unit in the LSTM layers is composed of a cell, an input gate, an output gate, and a forget gate. The cell remembers values over arbitrary time intervals, and the three gates regulate the flow of information into and out of the cell. The first gate is called the forget gate and it selects the part of the cell-state data to be removed. The second gate is the input gate, and it determines the data to be added to the cell state. Lastly, the output gate generates output data using the cell state. The output from the last cell is taken out, and a softmax function is applied to it. The final output is the probability of each class the model is predicting.

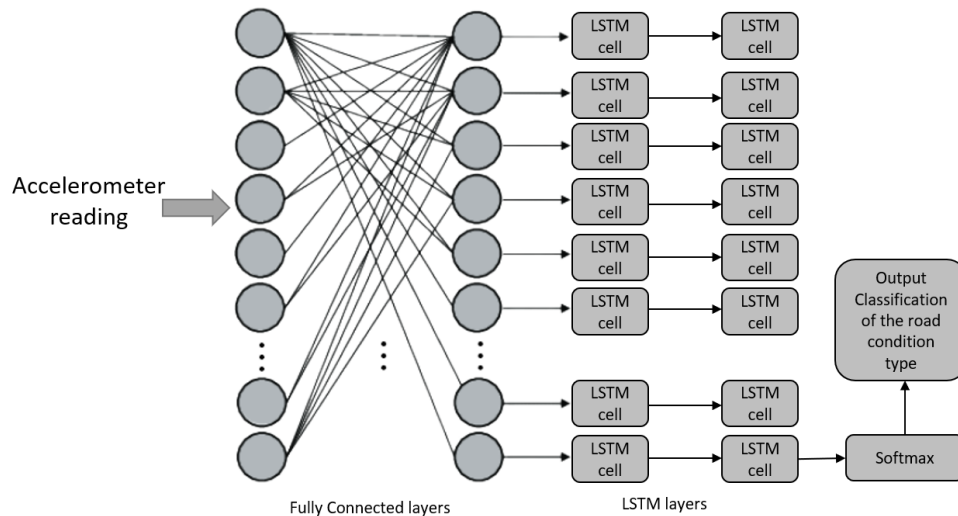


Figure 2 LSTM model structure for motion-based road damage detection.

3.3 Vision-Based Road-Damage Detection

Detecting road damage is one subdivision of an object-detection problem, which is a process that utilizes computer vision and image processing to deal with detecting instances of semantic objects of a certain class in digital images and videos. Among all the object detection methods, machine learning, especially deep convolutional neural networks, is the most prominent. Various network architectures, such as R-CNN, YOLO, R-FCN, and SSD, have been proposed. We selected YOLOv5 [28] and used transfer learning on it to achieve the detection of potholes and cracks. The training images were labeled by us to mark the areas of potholes and cracks, and the YOLOv5 model will catch the features of each road-damage type during training. The trained model will take the live video captured by the smartphone's camera as input and output bounding boxes and labels at the predicted areas of potholes and cracks in each frame.

In YOLOv5, the Leaky ReLU activation function is used in the middle/hidden layers, and the sigmoid activation function is used in the final detection layer. The default optimization function for training

is SGD. YOLO works by first dividing the image into grids, and, for each grid cell, it simultaneously produces bounding boxes, confidence, and class probability. The model then aggregates those results to make the final bounding box output and classification. This architecture is known for both its performance and efficiency. Over the years, there have been multiple iterations and improvements over the original YOLO architectures. This latest iteration, YOLOv5, utilizes cross stage partial network (CSPNet) as the model backbone and path aggregation network (PANet) as the neck for feature aggregation. These improvements have led to better feature extraction, and a significant boost in the mean averaged precision score.

3.4 Cloud-Based Collaborative Fusion

Cloud-based fusion is needed for the proposed method due to two reasons: Firstly, one vehicle cannot easily cover the entire road network, or even all road damage in one pavement section, during daily driving. However, the chances of covering all road damage and the entire road network increase if enough vehicles run the detection algorithm. A server on the cloud is required to gather this multi-vehicle data. Secondly, both vision-based and motion-based detection methods will unavoidably report false positives and false negatives, and such misleading information can be filtered out using cloud-based fusion. The type of damage and the GPS coordinates at which the damage is detected are used for the cloud-based collaborative fusion. The damage report is stored in a relational database in the cloud. The relational database has GPS coordinates (i.e., longitude and latitude), type of damage, and confidence as its attributes. Considering there will be multiple entries for the same location and reporting damage, there is a need to consolidate this reported damage, and this need is fulfilled by the cloud-based fusion. The complete architecture of this approach is shown in Figure 3. Amazon web services (AWS) Lambda function is used with the AWS API gateway for posting the data to the relational database. The raw data in the database is optimized using clustering techniques to populate another database with optimized data. Using the AWS Lambda function and AWS API gateway, this data is posted to the website.

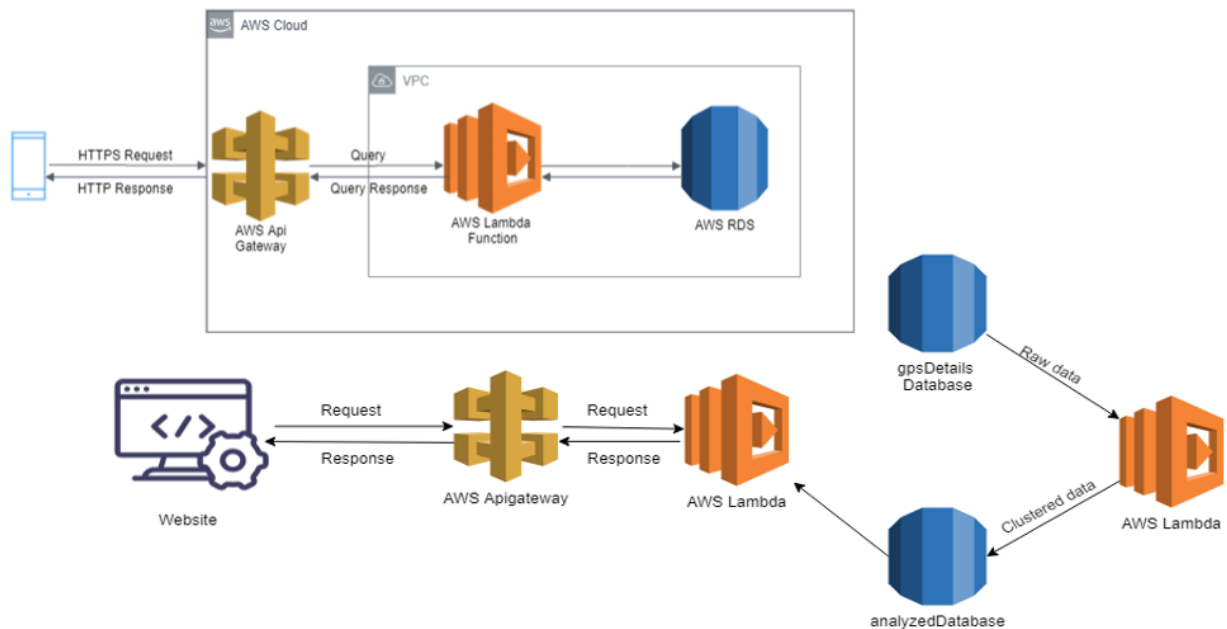


Figure 3 Architecture of cloud-based collaborative fusion.

Figure 4 gives a brief description of how the optimization of data is carried out on the reported damage. Different reports made at the same location might not have the exact same GPS coordinates, resulting in the database having a huge amount of duplicated damage reported for the same location.

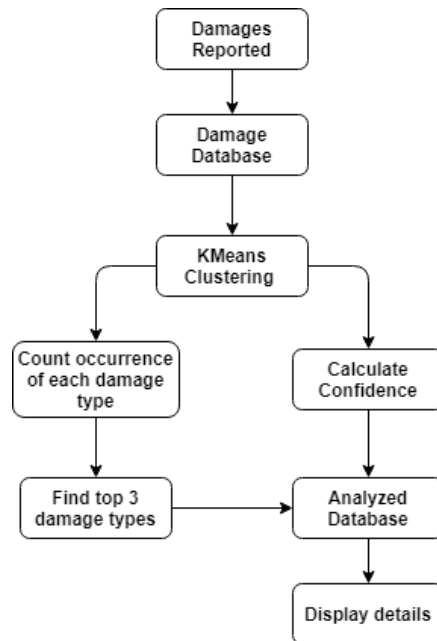


Figure 4 Cloud-based collaborative fusion flowchart.

With the need to consolidate these reports, an unsupervised machine-learning algorithm, k-means clustering, is implemented on the GPS coordinates. K-means is a numerical, unsupervised, non-deterministic, and iterative method [29]. Its efficiency and effectiveness in clustering have been proven in many practical applications. The algorithm starts with randomly selected k centers, where the value of k is pre-selected [30]. The algorithm then assigns each data point to its nearest center. The nearest center is generally calculated by using the Euclidean distance. This creates the initial k clusters [31]. After all the data points are assigned to their respective nearest center, the algorithm recalculates the centers. The new centers are calculated by averaging the data points that have been assigned to the initial centers. These recalculated centers are the new centers and the algorithm goes through all the data points again to assign them to the new nearest center. This recalculation and reassignment are repeated until the criterion function becomes minimum or the algorithm has looped for a predefined number of times.

In our approach, we use only the longitudes and latitudes of the reports from the data points for clustering. We start the implementation with an initial value of k as 1 and perform k-means clustering on the entire data set. Within clusters sum of squares (WCSS) is the measure of the average squared distance of all points within the clusters to the centroid of the cluster and is calculated for all k clusters; then, the average WCSS is calculated from the k WCSSs. Next, we increase k to $k + 1$ and repeat the process, until the difference between the average WCSS of the current $k + 1$ clusters and that of the previous k clusters is less than or equal to 0.001. Once that condition is reached, the number k is the optimum number of clusters that can be created given the current data set. The k centroids generated by applying k-means clustering are the longitude and latitude values representing the data points assigned to that particular cluster. The data points present in each cluster have a very high probability of being at the very same location or within a very small distance from one another.

We consolidated the entire data set of n rows into k clusters where k is less than n . For the k clusters to truly represent the entire data set, we need to analyze the data points in each cluster so that, along with the centroids representing the location of the cluster, the damage type and confidence can also be represented by the cluster. In order to achieve this, we segregated the data points according to the clusters they were assigned to. Each cluster and the associated data points were then visited to find the type of damage reported, and the occurrence of each type of damage in the

cluster was counted. The top three most occurring types of road damage within a cluster are the most possible or significant ones for that particular cluster. Then, we could represent the cluster location with the centroid values of that cluster and the most significant damage types by the three most occurring damage types, and we need to find the confidence for each of these damage types. This was performed by calculating the sum of confidence of each corresponding damage type and dividing it by the total number of occurrences. The process of finding the types and confidence was performed on all the k clusters. We, further, stored the analyzed information of the clusters in a cloud database. The database consists of the longitude and latitude, the centroid of the cluster, and the corresponding top three types of damage and their confidences. Additionally, we stored the cluster ID to distinguish different clusters in the database. This workflow is explained above, as shown in Figure 5.

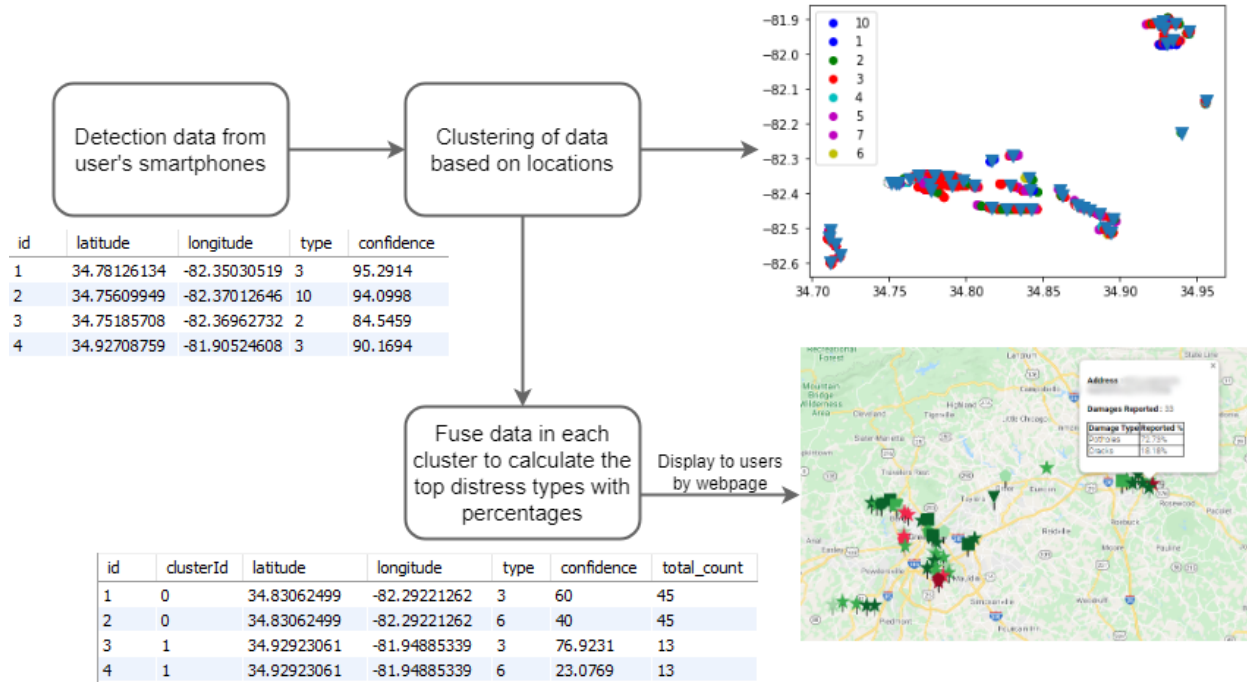


Figure 5 Back-end process of cloud-based collaborative fusion

3.5 Road-Damage Severity Estimation

Apart from utilizing vehicle-motion data to detect the existence of road damage, we also utilize motion data to estimate the severity of road damage to obtain comprehensive monitoring of road damage.

3.5.1 Method design and data acquisition

Generally speaking, the most direct impact from road damage is a vehicle bump, and a vehicle's bump can be reflected by vehicle vertical acceleration. On the other hand, vehicle vertical-acceleration change is mostly caused by vertical road-profile change. Therefore, it is important to detect if there exists any relationship between vehicle vertical acceleration and road-damage depth.

To inspect whether there is a reasonable relationship between vertical acceleration and road damage depth, we collected historical data on vehicle vertical acceleration on different road-damage sections. For each road-damage location, data from multiple runs were collected. In addition, we measured the length, width, and depth of road damage using tape measures. Once all data samples were collected, the next step was to find out the usage metrics. In this case, we used the different combinations of maximum value, average value, and variance to investigate what

leads to a reasonable prediction of road-damage severity levels. Take potholes as an example, assume we have M potholes of different sizes, and, for each pothole, collect acceleration data of M separate runs. Firstly, we find the average and maximum value of each run, which lead to an M maximum value and M average value. For both M maximum value and M average value, we then further calculate three statistics: their maximum value, average value, and variance, so that each pothole has six values to pair with damage depth. The chart is shown in Figure 6.




Figure 6 Data visualization of different road-damage estimation parameters.

Among all the combinations, the maximum of all the maximum acceleration data depicts a close-to-linear relationship to the depth of road damage. Thus, it is possible that the worst-case vehicle vertical acceleration can be used to estimate the road-damage severity. An example of a linear fit of data collected from 15 locations is shown in Figure 7.




Figure 7 Road-damage severity estimation based on acceleration data in the worst-case scenario.

3.5.2 Cloud-Based Road-Damage Severity Estimation

Taking the line fit result as a reference, the road-damage severity (depth, in this case) can be found from the worst-case vehicle vertical acceleration, which needs to be obtained from multiple different acceleration trajectories collected from the same location. Such different trajectories are generated

by different vehicles running over the road damage at different speeds or different approaching angles. A single trajectory from one vehicle cannot be the basis for a valid prediction since the vehicle may run over the road damage at different speeds, which results in different acceleration trajectories, or even not run over the road damage at all. Consequently, the aforementioned estimation method, which is based on the data of multiple vertical-acceleration trajectories, can only be achieved using cloud-based architecture so that the algorithm can have access to data from multiple vehicles. The cloud-based road-damage severity estimation process is shown in Figure 8. The process leverages the cloud server described in the previous section, and the vertical-acceleration trajectories of a vehicle will be sent to the cloud together with the prediction result. Data collected from the same location will be grouped based on the location using the k-means method. The maximum vertical acceleration of all available data is used to estimate the damage severity on that road section; thus, the estimation reliability increases as the size of the dataset increases.




Figure 8 Cloud-based road-damage severity estimation process

CHAPTER 4

Results

In this section, the data-collection process and experimental results will be presented in detail.

4.1 Experimental Setup and Data Collection

To conduct the experiment, we developed a dedicated data-collection app, which will be introduced in Section 4.1.1. The composition of the data collected using the app, as well as the data that are available from the existing data set, are presented in Section 4.1.2.

4.1.1 Setup and application development

This section includes a data-collection app [32] and explains how the data were collected using a smartphone from the sensors in smartphones. An existing android app was modified, which records the data from sensors—accelerometer, magnetometer, gyroscope, and GPS—along with the video of the roads where the data was being recorded, simultaneously. The UI of the app is as shown in Figure 9. The app gives two outputs—a video file and a .csv file that contains the sensor readings. During both data collection and real-time detection, the smartphone with this app is mounted onto a windshield, as shown in Figure 10.




Figure 9 Screen view of the data-collection app



Figure 10 Data-collection example

4.1.2 Data collection

The smartphone mounted on the windshield records the sensor data. The data used for this experiment was collected over various roads of Greenville, Spartanburg, Clemson, and Columbia area in South Carolina, USA. The data collection was performed using multiple cars. During data collection, the speed and placement of the vehicles when they pass over the road damage were also taken into consideration to produce as much diverse data as possible. The vehicles do not need to maintain constant speeds during data collection, but the final data set should contain data that covers almost the entire common speed range, such that the trained machine-learning models are free from the influence of vehicle speed, and the severity-estimation algorithm can have enough data to calculate the necessary statistics. The different types of road damage covered during data collection were bumps, construction joints, cracks, spalls, and normal road surfaces. The details of the collected 1378 data points can be seen in Table 1.

Table 1: Count of data across road-damage types.

Class	Count	Class	Count
Bump	39	Construction Joint	156
Crack	721	Undamaged	310
Spall	152		

The smartphone optimizes the use of sensors, changing the frequency of data recording; hence, not all the collected data are usable. The usable data points are the ones with a sampling frequency of 60 Hz or higher. The details of the usable motion data for detection are presented in Table 2.

Table 2: Count of acceleration data across classes.

Class	Count
Damaged	378
Undamaged	310

The images from [33], one of the largest road-damage data sets, pothole data set [34], and images extracted from the video frames from the data recording were used for the vision-based YOLOv5 model as well. After careful filtering and processing, 1000 images were retained. Data-augmentation techniques, such as flipping in all four directions, rotating by 90° left and right, and 20% crop, were implemented on the original images. The images were increased to 3000 after the application of the augmentation techniques. This process introduced variations and increases the size of the data set. Table 3 shows the breakup of images into classes—potholes and cracks—after the augmentation techniques.

Table 3: Count of images across classes after augmentation

Class	Count
Potholes	893
Cracks	767
Total Annotations	1660
Total Images after Augmentation	3000

4.2 Experimental Result and Analysis

4.2.1 Results of Vision-Based Road-Surface Detection

YOLOv5 was trained extensively for the road-damage types, including potholes and cracks, for 200 epochs. The model performed well on the validation data set with an accuracy of 87.5% mean average precision (MAP). The MAP improvement over the epochs is shown in Figure 11. The MAP gradually improved as the training progressed initially, and eventually almost stopped changing,

which indicates that the model was trained successfully. Figure 12 shows the training and validation loss over the epochs. Both the training loss and the validation loss dropped to close to zero at the end of the training, meaning that the model was not encountering the overfitting issue.

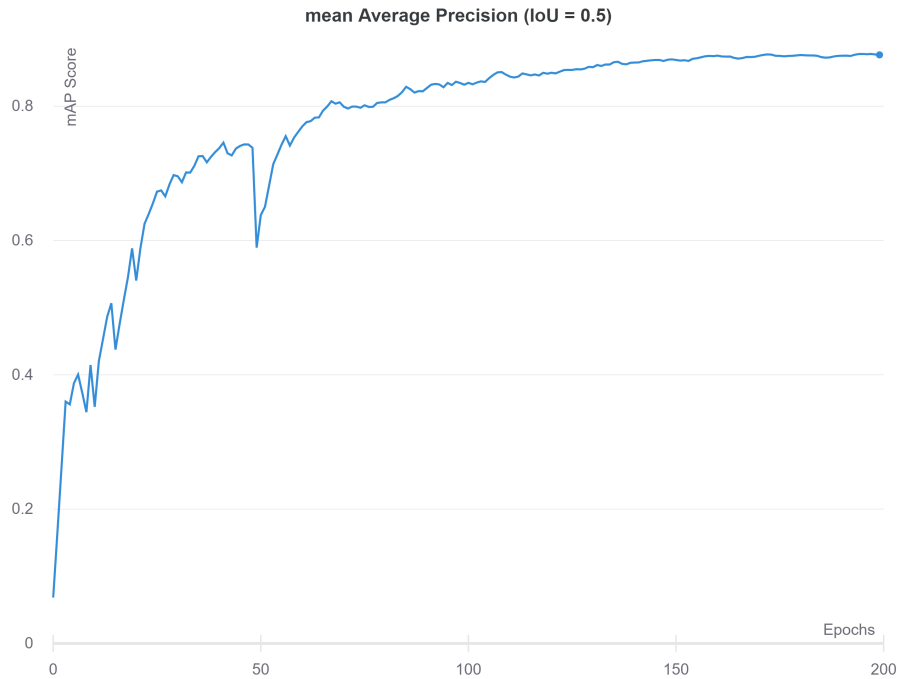


Figure 11 YOLOv5 mean average precision over epochs

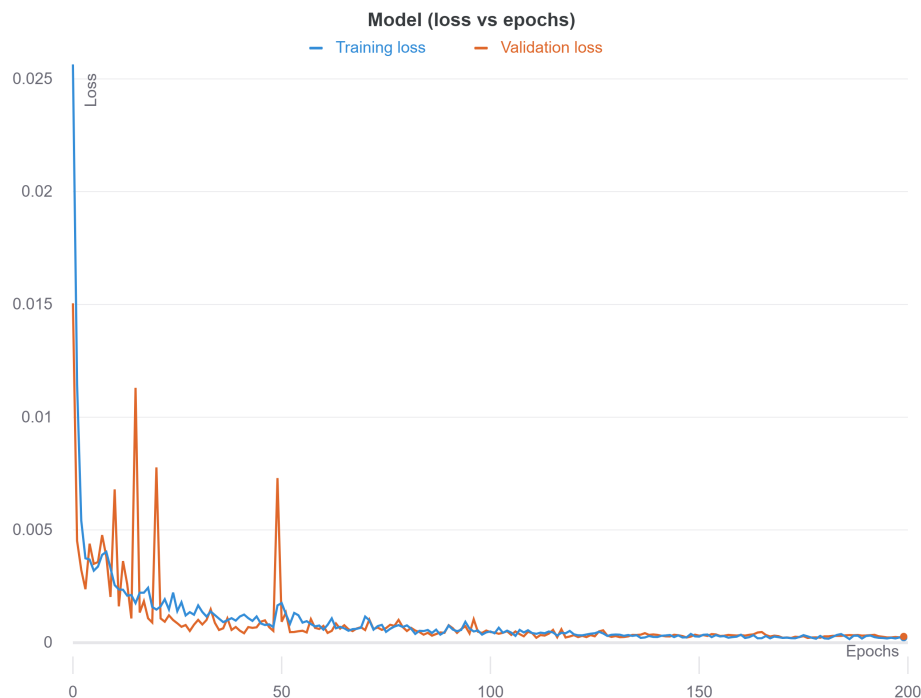


Figure 12 YOLOv5 loss over epochs on the training and validation data set.

Figure 13 and Figure 14 show the confusion matrix of the inferences on the training and testing data set, respectively. This matrix compares the actual target values with those predicted by our YOLOv5 model. The rows represent the predicted values of the target variable.

The column represents the ground truth of the target variable. Table 4 and Table 5 provide the summary of the trained model accuracy on the training and testing data set, respectively. The performance numbers were obtained when the threshold of intersection over union (IOU) was set to 0.5 for successful detection. The results show that the trained model never confuses cracks and potholes. On the training dataset, the model did not miss hardly any detection. On the testing dataset, the model did not miss hardly any potholes but failed to detect 22% of the cracks. On both the training and the testing datasets, the model generated some false positives when there was no road damage. The calculated F1 score is 0.853, which is better than the performance of the networks for the same purpose in [18]–[23]. Overall, the results show that the trained YOLO model is especially good at telling cracks from potholes, but can sometimes make mistakes in distinguishing undamaged roads from damage.

Table 4: Accuracy results of training data sets

Class	Annotations	Precision	Recall	mAP@0.5IOU
Potholes	1826	0.898	0.999	0.997
Cracks	1621	0.999	0.984	0.996
Total	3447	0.994	0.991	0.997

Table 5: Accuracy results of testing data sets

Class	Annotations	Precision	Recall	mAP@0.5IOU
Potholes	779	0.932	0.893	0.937
Cracks	660	0.888	0.730	0.813
Total	1439	0.910	0.812	0.875




Figure 13 Confusion matrix—YOLOv5 model on the training data set.




Figure 14 Confusion matrix—YOLOv5 model on the testing data set.

The precision vs. recall curve for the two classes, potholes and cracks, and the overall detection is shown in Figure 15. The model was able to give 81% MAP on cracks, and 91% MAP on potholes, after training for 200 epochs. It can be seen that all three curves are very close to the top right corner, which indicates that the trained model can be an effective predictor. Figure 16 shows the model inference on some of the images from the validation data set. The two classes of potholes and cracks were inferred well by our model.




Figure 15 YOLOv5 precision-recall curve.

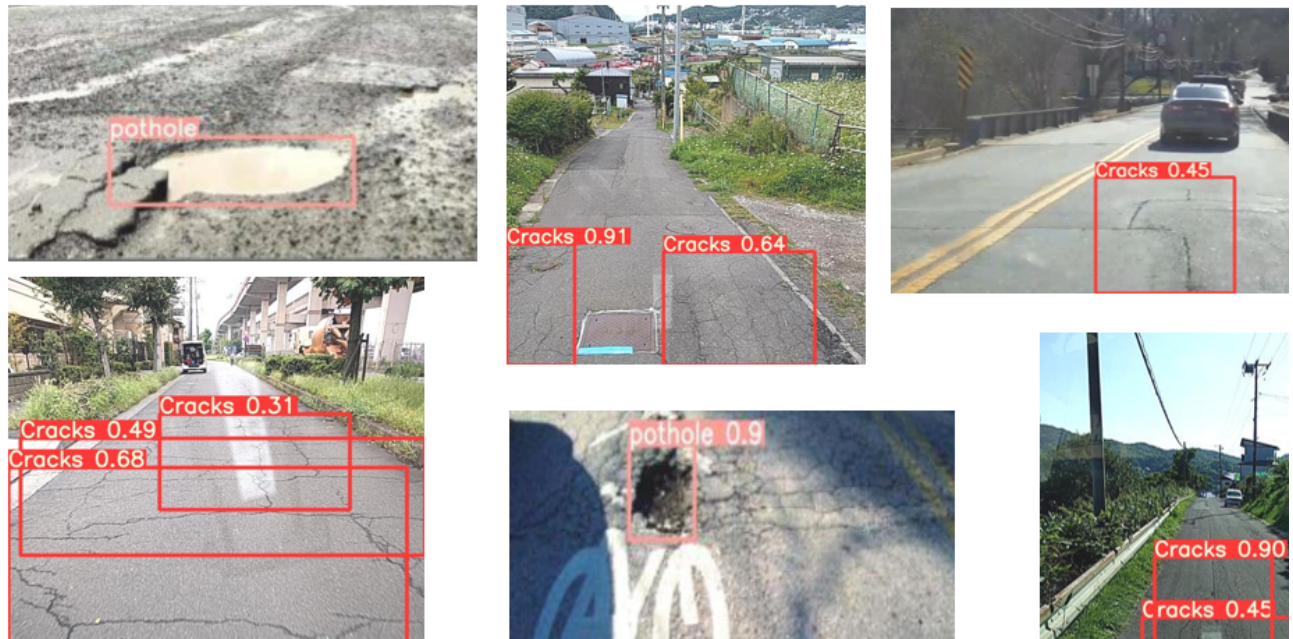


Figure 16 YOLOv5 inference on validation images.

4.2.2 Results of Motion-Based Road-Damage Detection

The LSTM model was trained on a Palmetto cluster on GPU nodes. The LSTM model was tuned well to obtain the necessary results. Table 6 gives the details on the tunable parameters used for the training. Figure 17 shows the prediction accuracy over epochs during training and Figure 18 shows how the loss reduced over epochs. It can be seen that the training was stopped when the model's performance almost stopped increasing on both the testing dataset and validation dataset and started to show differences between them. This indicates that the model was trained well before an overfitting issue started to appear. The trained LSTM model gives an accuracy of 94% in classifying the damage and normal road on the testing data set, performing better than models for similar purposes in [7]–[12]

Table 6: Tuning parameters of LSTM

Parameter Name	Value Used
Optimizaer	Adam
Loss regularization	L2(L2 loss used is 0.015)
Learning rate	0.005
Batch size	64




Figure 17 YOLOv5 inference on validation images.

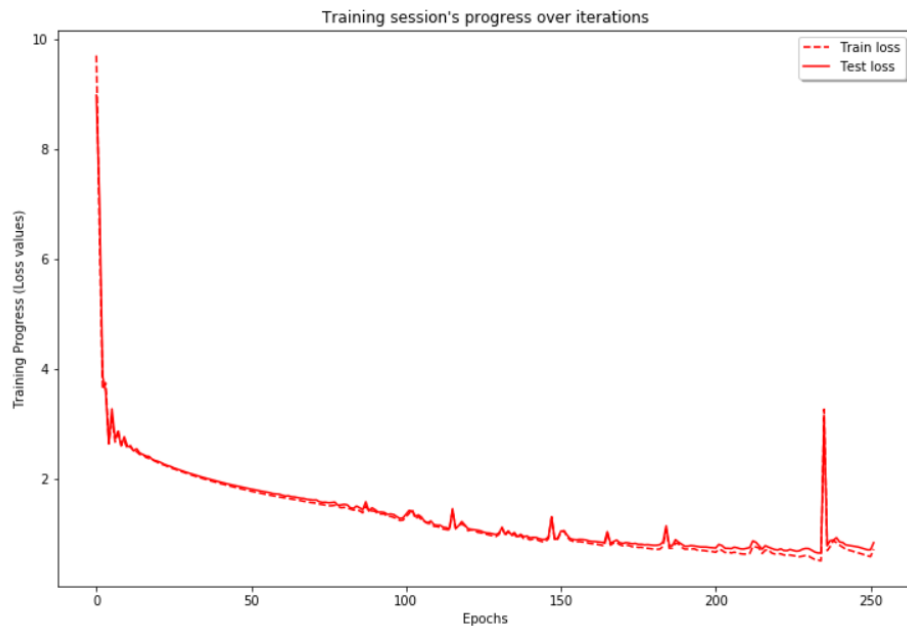


Figure 18 LSTM classification loss over epochs.

Figure 19 and Figure 20 show the confusion matrix of the model on training and testing data sets, respectively. The trained LSTM model is able to achieve a 99% precision for detecting undamaged roads and 96% precision for detecting damaged roads on the training dataset, and 96% precision for detecting undamaged roads and 90% precision for detecting damaged roads on the testing dataset. The model generated very few false positives on both datasets. Overall, the LSTM model performs very well in telling damaged road surfaces from undamaged ones and can be a very good compensation for the trained YOLO model.

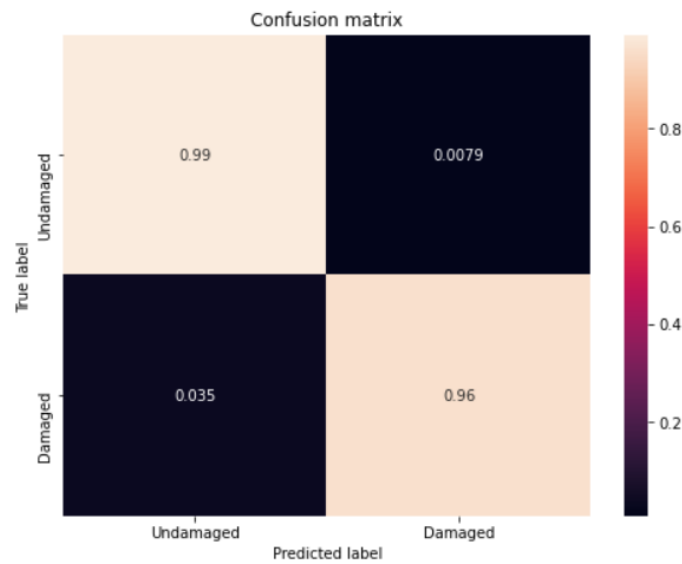


Figure 19 Confusion matrix—LSTM model on the training data set.




Figure 20 Confusion matrix—LSTM model on the testing data set.

4.2.3 Results of Cloud-Based Collaborative Fusion

Figure 21 shows a screenshot of the web page displaying the results of the clustering with the top three damage types and their corresponding confidence with the address of the damage. The web page enables concerned authorities to view the road damage reported by the users with the help of the mobile application. The web page comprises a map on which the locations with road damage are marked with different types of markers. The markers are of different shapes to distinguish between the highest reported type of damage in that location. There are a total of five different shapes for markers for the current two types of damage and three other types of damage to be included in the future. The marker shape and related damage type are as follows:

- Star: pothole;
- Square: cracks;
- Others: reserved;

Further, there is the ability to display the severity of the damage at a particular location, and the markers have five different colors. The five colors range from a dark green shade displaying less

severe damage to a dark red color displaying severe damage. The web-page viewer can click on the markers to know more details about the damage, such as type, confidence level, and location details. For example, the location shown in the figure has 33 different damage reports with high confidence after clustering, and the most frequent damage type is potholes, which constitutes 24 out of the 33 reports, and, thus, the location most likely has potholes on the road surface. When comparing our cloud-based fusion solution with other methods, our method has the advantages of tolerating false predictions and low requirements for single-vehicle data.

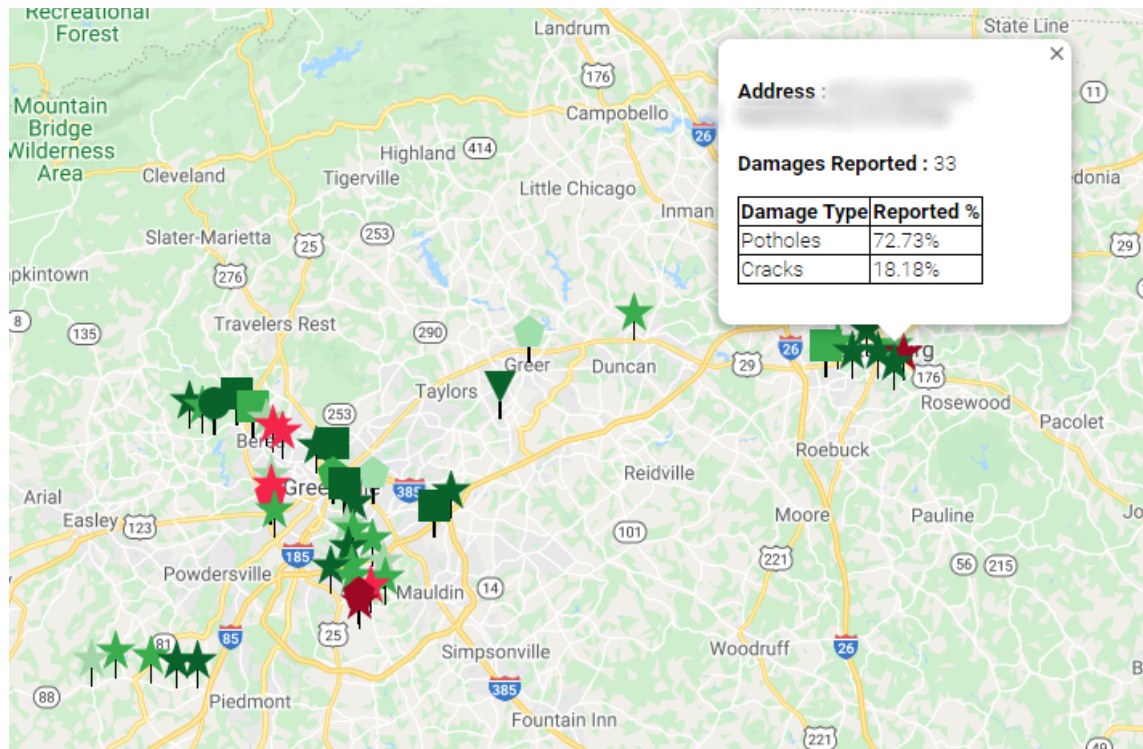


Figure 21 User interface web page.

4.2.4 Results of Cloud-Based Road-Damage Severity Estimation

In this section, we use data collected from 19 different damage locations to perform line fitting to build the mapping between vehicle vertical acceleration and road-damage depth. The data points used for line fitting are marked as blue points in Figure 22. Data from the other six damaged locations are used to test the accuracy of the road-damage estimation, which is marked as orange points in Figure 22. It can be seen that the testing data also match with the fitted line nicely.

To better demonstrate the accuracy of the road-damage severity estimation, an error analysis was also conducted. A detailed comparison between measured damage depth and estimated damage depth for the six testing locations can be found in Table 7. The average error is 1.79 cm based on the testing data, with a minimum error of 0.82 cm and a maximum error of 2.92 cm. Overall, the road-damage severity estimation method using the worst-case vertical acceleration and line-fitting technique can work effectively.




Figure 22 Damage-severity estimation validation.

Table 7: Road-damage-severity estimation error

Estimated data (cm)	8.82	11.88	8.7	8.41	7.76	6.92
Measured data (cm)	8	9	10	11	8	4
Error (cm)	0.82	2.88	1.3	2.59	0.24	2.92

CHAPTER 5

Conclusions

This project focuses on the problem of monitoring road-surface conditions for detecting possible road damage. A deep-learning and cloud-based collaborative method is proposed. The method utilizes both vehicle-motion data and vision data collected by cellphones to achieve reliable and accurate detection while being cost-effective. Deep-learning-based techniques including YOLOv5 and LSTM were trained for this task. The cloud fuses detecting results from various vehicles to generate a map with credible road-damage information for drivers. The performance of the proposed method was proven in real-world experiments. However, the work can still be improved in the following aspects: The current dataset we use to train the deep-learning models can be expanded to be more comprehensive. In the future, we will collect data from more weather and lighting conditions and vehicle types. The current severity-estimation method is mainly based on motion data. In our future work, we will combine the current estimation method with information from vision data to achieve the estimation of the 3-D dimensions of road damage. In addition, the current method does not output the pavement-quality index, which is a standard index for evaluating road-surface conditions. Our future work will create an algorithm to estimate this from vehicle vertical acceleration and speed. The current cloud-based fusion algorithm has its limitations as well. Firstly, it is not very time-efficient and may have some duplicated predictions in the cloud. In the future, we will investigate more efficient fusion algorithms that can better process the duplicated and redundant information from the cloud. Secondly, the current fusion algorithm directly fuses the prediction results from cellphones, which does not fully utilize the computational power of the cloud server. Our future work will offload computational tasks from cellphones and run more complex and precise deep-learning models on the cloud to further improve prediction accuracy. Lastly, the current cloud is centralized, which is not ideal for a large area, since the communication delay increases and fusion speed drops as the map grows. Our future work will improve the cloud structure to a distributed one to accommodate large areas.

REFERENCES

- [1] S. Singh, "Critical reasons for crashes investigated in the national motor vehicle crash causation survey," 2015.
- [2] P. Tang, D. Huber, and B. Akinci, "Characterization of Laser Scanners and Algorithms for Detecting Flatness Defects on Concrete Surfaces," *Journal of Computing in Civil Engineering*, vol. 25, no. 1, pp. 31–42, Jan. 2011.
- [3] A. González, E. J. O'Brien, Y.-Y. Li, and K. Cashell, "The use of vehicle acceleration measurements to estimate road roughness," *Vehicle system dynamics*, vol. 46, no. 6, pp. 483–499, 2008.
- [4] A. Mednis, A. Elsts, and L. Selavo, "Embedded solution for road condition monitoring using vehicular sensor networks," in *2012 6th International Conference on Application of Information and Communication Technologies (AICT)*, 2012, pp. 1–5.
- [5] K. De Zoysa, C. Keppitiyagama, G. P. Seneviratne, and W. Shihaan, "A public transport system based sensor network for road surface condition monitoring," in *2007 workshop on Networked systems for developing regions*, 2007, pp. 1–6.
- [6] F. E. Gunawan, B. Soewito, and others, "A vibratory-based method for road damage classification," in *2015 International Seminar on Intelligent Technology and Its Applications (ISITIA)*, 2015, pp. 1–4.
- [7] A. Mednis, G. Strazdins, R. Zviedris, G. Kanonirs, and L. Selavo, "Real time pothole detection using android smartphones with accelerometers," in *2011 International conference on distributed computing in sensor systems and workshops (DCOSS)*, 2011, pp. 1–6.
- [8] G. Sebestyen, D. Muresan, and A. Hangan, "Road quality evaluation with mobile devices," in *2015 16th International Carpathian Control Conference (ICCC)*, 2015, pp. 458–464.
- [9] R. Bhoraskar, N. Vankadhara, B. Raman, and P. Kulkarni, "Wolverine: Traffic and road condition estimation using smartphone sensors," in *2012 4th International Conference on Communication Systems and Networks, COMSNETS 2012*, 2012.
- [10] M. Perttunen, O. Mazhelis, F. Cong, M. Kauppila, T. Leppänen, J. Kantola, J. Collin, S. Pirttikangas, J. Haverinen, T. Ristaniemi, and others, "Distributed road surface condition monitoring using mobile phones," in *International conference on ubiquitous intelligence and computing*, 2011, pp. 64–78.
- [11] B. Varona, A. Monteserin, and A. Teyseyre, "A deep learning approach to automatic road surface monitoring and pothole detection," *Personal and Ubiquitous Computing*, vol. 24, no. 4, pp. 519–534, 2020.
- [12] J. Menegazzo and A. von Wangenheim, "Multi-Contextual and Multi-Aspect Analysis for Road Surface Type Classification Through Inertial Sensors and Deep Learning," in *2020 X Brazilian Symposium on Computing Systems Engineering (SBESC)*, 2020, pp. 1–8.
- [13] A. Chatterjee and Y.-C. Tsai, "Training and testing of smartphone-based pavement condition estimation models using 3d pavement data," *Journal of Computing in Civil Engineering*, vol. 34, no. 6, p. 4020043, 2020.
- [14] Y. Li, C. Liu, Y. Shen, J. Cao, S. Yu, and Y. Du, "RoadID: A Dedicated Deep Convolutional Neural Network for Multipavement Distress Detection," *Journal of Transportation Engineering, Part B: Pavements*, vol. 147, no. 4, p. 4021057, 2021.
- [15] Y.-C. Tsai, V. Kaul, and R. M. Mersereau, "Critical assessment of pavement distress segmentation methods," *Journal of transportation engineering*, vol. 136, no. 1, pp. 11–19, 2010.
- [16] D. Llopis-Castelló, R. Paredes, M. Parreño-Lara, T. Garcí'a-Segura, and E. Pellicer, "Automatic classification and quantification of basic distresses on urban flexible pavement through convolutional neural networks," *Journal of Transportation Engineering, Part B: Pavements*, vol. 147, no. 4, p. 4021063, 2021.
- [17] M. Alipour, D. K. Harris, and G. R. Miller, "Robust pixel-level crack detection using deep fully convolutional neural networks," *Journal of Computing in Civil Engineering*, vol. 33, no. 6, p. 4019040, 2019.
- [18] W. Wang, B. Wu, S. Yang, and Z. Wang, "Road damage detection and classification with faster R-CNN," in *2018 IEEE international conference on big data (Big data)*, 2018, pp. 5220–5223.
- [19] K. Doshi and Y. Yilmaz, "Road damage detection using deep ensemble learning," in *2020 IEEE International Conference on Big Data (Big Data)*, 2020, pp. 5540–5544.
- [20] S. Naddaf-Sh, M.-M. Naddaf-Sh, A. R. Kashani, and H. Zargarzadeh, "An efficient and scalable deep learning approach for road damage detection," in *2020 IEEE International Conference on Big Data (Big Data)*, 2020, pp. 5602–5608.

- [21] V. Hegde, D. Trivedi, A. Alfarrarjeh, A. Deepak, S. H. Kim, and C. Shahabi, "Yet another deep learning approach for road damage detection using ensemble learning," in *2020 IEEE International Conference on Big Data (Big Data)*, 2020, pp. 5553–5558.
- [22] V. Mandal, A. R. Mussah, and Y. Adu-Gyamfi, "Deep learning frameworks for pavement distress classification: A comparative analysis," in *2020 IEEE International Conference on Big Data (Big Data)*, 2020, pp. 5577–5583.
- [23] T. Hascoet, Y. Zhang, A. Persch, R. Takashima, T. Takiguchi, and Y. Ariki, "Fasterrcnn monitoring of road damages: Competition and deployment," in *2020 IEEE International Conference on Big Data (Big Data)*, 2020, pp. 5545–5552.
- [24] V. Pham, C. Pham, and T. Dang, "Road damage detection and classification with detectron2 and faster r-cnn," in *2020 IEEE International Conference on Big Data (Big Data)*, 2020, pp. 5592–5601.
- [25] M. A. Ameddah, B. Das, and J. Almhana, "Cloud-assisted real-time road condition monitoring system for vehicles," in *2018 IEEE Global Communications Conference (GLOBECOM)*, 2018, pp. 1–6.
- [26] A. Ghose, P. Biswas, C. Bhaumik, M. Sharma, A. Pal, and A. Jha, "Road condition monitoring and alert application: Using in-vehicle smartphone as internet-connected sensor," in *2012 IEEE international conference on pervasive computing and communications workshops*, 2012, pp. 489–491.
- [27] Y. Yuan, M. S. Islam, Y. Yuan, S. Wang, T. Baker, and L. M. Kolbe, "EcRD: Edge-cloud computing framework for smart road damage detection and warning," *IEEE Internet of Things Journal*, vol. 8, no. 16, pp. 12734–12747, 2020.
- [28] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [29] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.
- [30] S. Na, L. Xumin, and G. Yong, "Research on k-means clustering algorithm: An improved k-means clustering algorithm," in *2010 Third International Symposium on intelligent information technology and security informatics*, 2010, pp. 63–67.
- [31] A. K. Jain, M. N. Murty, and P. J. Flynn, "Data clustering: a review," *ACM computing surveys (CSUR)*, vol. 31, no. 3, pp. 264–323, 1999.
- [32] K. A. A. Nazeer and M. P. Sebastian, "Improving the Accuracy and Efficiency of the k-means Clustering Algorithm," in *The World Congress on Engineering*, 2009, vol. 1, pp. 1–3.
- [33] R. Karthik, E. Cuclurciello, dujianyi, and D. Rosson, "e-lab video sensor app for data recording." 2017.
- [34] H. Maeda, T. Kashiyaama, Y. Sekimoto, T. Seto, and H. Omata, "Generative adversarial network for road damage detection," *Computer-Aided Civil and Infrastructure Engineering*, vol. 36, no. 1, pp. 47–60, 2021.