PLEASE LEAVE THIS BOX IN PLACE AND DO NOT TYPE ANYTHING WITHIN THE BOX

Abstract

Deteriorating roads pose safety risks to road users and can cause costly damage to the vehicles. The severity of hazards caused by road defects can range from minor to severe. These hazards can be minimised by the timely detection of road defects. Technological advancements have led to traditional inspection methods such as visual inspection being replaced by more advanced methods such as deep learning techniques for autonomous road defect detection. However, one of the major challenges faced by deep learning techniques is the requirement of significant amounts of training data. The acquisition of large amounts of data is rather costly due to equipment, vehicle fuel and data storage expenses. Additionally, integrating deep learning models for road defect detection with existing international codes and standards remains a challenge. This short paper presents a quicker and more efficient data acquisition method for acquiring data to train a deep learning road defect detection model using transfer learning. The model was also designed to allow for easy integration with the UK highway inspection manual. The model demonstrated good performance, achieving precision, recall and mAP values of 89.5%, 81.6% and 84.6%, respectively.

Keywords: Road inspection, road defect, object detection, google street view, deep learning, transfer learning.

1 Introduction

Global population and economic growth have inevitably led to an increase in the traffic volumes. This increase has subsequently led to an increase in the traffic loads [1] exerted on roads, resulting in a diverse range of road defects. The severity of these defects can vary from minor inconveniences to severe hazards which threaten human safety and vehicle integrity. Timely detection of defects is essential to ensure appropriate planning of relevant maintenance and repair strategies [2].

Inspections to determine the condition of roads have predominantly been carried out using manual methods such visual inspection. Although road defects can be identified accurately through manual visual inspection, this process tends to be tedious and subjective [1]. Technology advancements have seen the use of digital cameras and image processing methods being used in road defect detection. Most recently, road defect detection has shifted towards the use of deep learning techniques due to their quicker, less subjective capabilities. Majority of the existing studies have focussed on the detection of cracks in roads [3–7] using RGB images. Chun et al. [4] performed classification of 6 classes: crack, non-crack, road markings with cracks, road markings without cracks, facilities with cracks and facilities without cracks. Radapoulou [2] detected longitudinal cracks, transverse cracks, patches and potholes. Chen et al. [1] detected transverse cracks, longitudinal cracks, joints/patches, potholes, manholes and shadows, road markings and oil stains using thermal images.

While majority of existing methods are effective in road defect detection, they are faced with a few limitations. Firstly, collecting images for road defect datasets is expensive due to the extensive data collection, storage, and processing involved. This is because inspectors might find themselves driving for many miles in search of various types of defects and as result incurring high fuel costs. In most cases, roads are not so excessively damaged in one area that all defect types can be found with during a short drive. This results in data storage and processing limitations arising such as large amounts of video footage and tedious data preprocessing to sift through thousands of mostly defect free images just to get to the images with defects. This often results in a class imbalance, where there is a disproportionate representation of certain defect types in the dataset. For instance, Opara et al. [8] manually collected images using a Road Space Information Management (RIM) vehicle. The final dataset consisted of 1035, 1676, 672, 1968 and 11 images of transverse, longitudinal, alligator, no crack and pothole images, respectively. The authors attributed the class imbalance in the study to the high efficiency of Japanese authorities in promptly repairing potholes, making such imbalances inevitable. Some studies [9–11] have used drones to overcome the issue of long drives in search for defects however, these are often limited by flight range and battery life.

Secondly, the identification of distress types often relies on the intuitive understanding and recognition of prevalent defects. To the authors' best knowledge, none of the existing DL models have been created to support and be used in conjunction with international codes and standards for road defect detection. It is essential for DL models to be designed for easy integration with existing codes, this will facilitate their deployment beyond the laboratory to field applications. To overcome the limitations of current DL methods for automated road defect detection, this paper presents a method that utilises data from Google Maps Street view to train a pre-trained object detection deep learning model by transfer learning.

2 Methods

2.1 Data collection and processing

Images were collected by manually walking through Google Street View and capturing screenshots using the snipping tool in windows. The images were collected from a number of streets in Cardiff, Swansea and Nottingham, UK. The images collected from Google Street View were pre-processed and manually labelled in Roboflow [12]. A total of 150 images were collected for initial training of the road defect detection model. These images underwent a several data augmentation steps to increase the dataset size. The augmentation of the images led to an improvement in the model's ability to generalise on new and unseen data which can be encountered when inspecting roads in different lighting, weather conditions and in the presence of noise and blurriness. After augmentation, the image dataset was increased in size to 660 images.

Four augmentation steps were implemented:

1. Rotation between -8% and +8%
2. Exposure adjustment between -10% and 10%
3. Blur up to 1.8px
4. Noise up to 1.0% of pixels

Three defects were identified in the images using bounding boxes, these made up the classes to be detected by the model. The type of road destresses labelled in the images were potholes, cracks, patches. Figure 1 shows an example of two images which were labelled using Roboflow. To allow for easy integration with the UK Highway inspection manual [14], these classes were named as found in the inspection manual. Potholes were referred to as CW01-Potholes and Cracks were referred to as CW08-Cracks. The "CW" indicates that these defects were detected on the carriage way. While the third defect was not listed in the UK highway inspection manual, it was included as a class to help the model to differentiate between potholes and patches. This is because these two classes can easily be mistaken for each other.

Figure 1: Labelled images showing potholes, cracks and patches.

2.2 Deep learning model

To train the model to detect the defects labelled in the images, transfer learning was utilised due to the small size of the dataset. The use of transfer learning helps overcome this problem and improves accuracy of results by repurposing already built models previously trained on larger datasets for different tasks. The object detection model YOLO-NAS [13], was employed for this task of road defect detection using transfer learning.

2.3 Performance evaluation

The performance of the DL model was evaluated using several evaluation metrics. These metrics were precision, recall and mAP. These metrics are defined as follows:

$$\text{Precision} = \frac{TP}{TP+FP} \tag{1}$$

$$\text{Recall} = \frac{TP}{TP+FN} \tag{2}$$

$$\text{mAP} = \frac{1}{N}\sum_{i=}^{N} AP_i \tag{3}$$

Where TP represents True Positives, TN represents True Negatives, FP represents False Positives, FN represents False Negatives, AP is Average Precision for class i and N is the number of classes.

Using the metrics presented above the model's performance in road defect detection task was evaluated and the results are presented in the next section.

3 Results

The DL model was trained for a total of 150 epochs over a duration of 3 hours. The model achieved a mAP of 84.6%, precision of 89.5% and a recall of 81.6%. The training graphs for recall, precision and mAP are illustrated in Figure 2. A mAP of 84.6% indicated that the model correctly detected defects across all the three classes 84.6% of the time. Figure 3 further illustrates the performance of the model across the three classes, with respect to precision. It can be seen that the model correctly predicted patches, 100% of the time. Cracks (CW08-Crack) were correctly predicted 85% of the time, while potholes (CW01-Potholes) were only predicted correctly 69% of the time. To improve the precision of the model, more images of the CW01-pothole and the CW08-Crack classes need to be added to the training dataset.

A recall of 81.6% was achieved by the model which means that the model was able to find 81.6% of the defects in the images in the dataset, leaving only 18.4% of defects undetected. Improving the recall score would reduce the likelihood of defects going undetected in the real-world inspections.

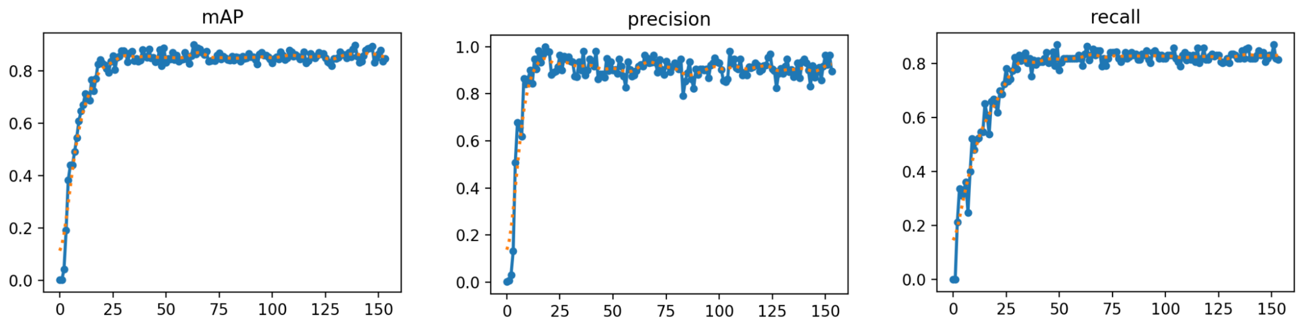

(a) (b) (c)

Figure 2: Training curves showing (a) mAP curve, (b) Precision curve and (c) recall curve.

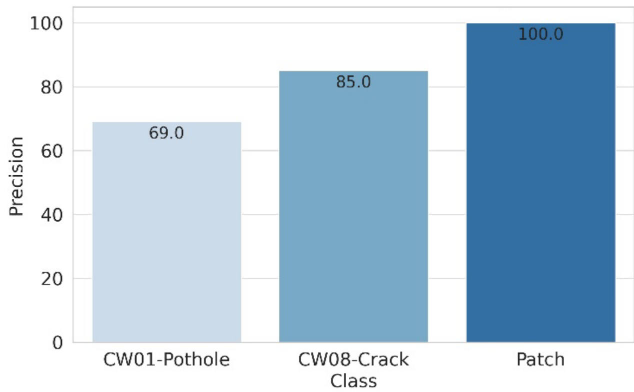

Figure 3: Precision comparison for the three classes of road defects.

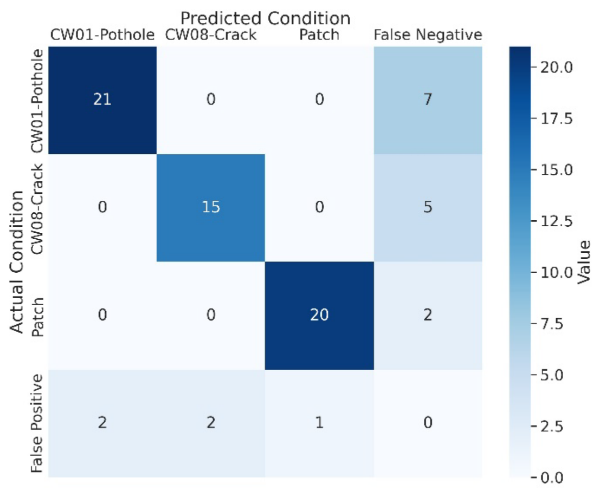

Figure 4: Confusion matrix showing results from the test dataset.

Figure 4 presents a confusion matrix displaying model’s performance on a 35 test images with 75 instances of the three defect classes. It can be seen that patches had the least number of FN and FP. While CW01-Potholes were on the other end of the spectrum, with the highest number of FN and FP.

Figure 5 illustrates a typical output from the model, which includes bounding boxes around detected defects, accompanied by the defect name and the model's confidence in its prediction. The defects in Figure 5(a) were detected with high confidence, while those in Figure 5(b) which were detected with lower confidence. This lower confidence in Figure 5(b) can be attributed to the small dataset size. To improve the confidence scores and overall performance of the model, the dataset will be increased in future work to allow the model to be trained on a variety of data. This will in turn lead to better generasilibilty of the model.

(a) (b)

Figure 5: Detected road defects in two scenarios.

4 Conclusions and Contributions

The short paper has presented a DL model for detecting road defects using transfer learning and google street view for collecting the training data. The following conclusions can be drawn from the paper:

- Utilising Google Street View for data collection has proven to be an effective and accurate method for training models from a desktop setting. This approach significantly reduces the time and costs associated with fuel, data storage, and expensive equipment needed for physical data collection.
- The transfer learning model was trained to identify defects in line with the UK Highways Inspection Manual, demonstrating its potential for practical field applications.

References

[1] C. Chen, S. Chandra, Y. Han, H. Seo, "Deep Learning-Based Thermal Image Analysis for road Defect Detection and Classification Considering Complex road Conditions", Remote Sensing, 14, 106, 2021. doi:10.3390/RS14010106.

[2] S.C. Radopoulou, I. Brilakis, "Automated Detection of Multiple road Defects", Journal of Computing in Civil Engineering, 31, 04016057, 2016. doi:10.1061/(ASCE)CP.1943-5487.0000623.
[3] Y. Shi, L. Cui, Z. Qi, F. Meng, Z. Chen, "Automatic Road Crack Detection Using Random Structured Forests", IEEE Transactions on Intelligent Transportation Systems, 17, 3434-3445, 2016. doi:10.1109/TITS.2016.2552248.
[4] P.J. Chun, T. Yamane, Y. Tsuzuki, "Automatic Detection of Cracks in Asphalt road Using Deep Learning to Overcome Weaknesses in Images and GIS Visualization", Applied Sciences, 11, 892, 2021. doi:10.3390/APP11030892.
[5] Q. Zou, Y. Cao, Q. Li, Q. Mao, S. Wang, "CrackTree: Automatic Crack Detection from road Images", Pattern Recognition Letters, 33, 227-238, 2012. doi:10.1016/J.PATREC.2011.11.004.
[6] F. Yang, L. Zhang, S. Yu, D. Prokhorov, X. Mei, H. Ling, "Feature Pyramid and Hierarchical Boosting Network for road Crack Detection", IEEE Transactions on Intelligent Transportation Systems, 21, 1525-1535, 2020. doi:10.1109/TITS.2019.2910595.
[7] Z. Fan, H. Lin, C. Li, J. Su, S. Bruno, G. Loprencipe, "Use of Parallel ResNet for High-Performance road Crack Detection and Measurement", Sustainability, 14, 1825, 2022. doi:10.3390/SU14031825.
[8] J.N. Opara, A.B.B. Thein, S. Izumi, H. Yasuhara, P.-J. Chun, "DEFECT DETECTION ON ASPHALT road BY DEEP LEARNING", GEOMATE Journal, 21, 87-94, 2021. doi:10.21660/2021.83.6153.
[9] M. Petkova, P. Maes, "Deploying Drones for Autonomous Detection of road Distress", 2016.
[10] Q. Qiu, D. Lau, "Real-Time Detection of Cracks in Tiled Sidewalks Using YOLO-Based Method Applied to Unmanned Aerial Vehicle (UAV) Images", Automation in Construction, 147, 104745, 2023. doi:10.1016/J.AUTCON.2023.104745.
[11] I. Katsamenis, N. Bakalos, E. Protopapadakis, E.E. Karolou, G. Kopsiaftis, A. Voulodimos, "Real Time Road Defect Monitoring from UAV Visual Data Sources", ACM International Conference Proceeding Series, 603-609, 2023. doi:10.1145/3594806.3596561.
[12] Roboflow Inc, "Roboflow Annotate: Label Faster Than Ever", Available online: https://roboflow.com/annotate (accessed on 3 September 2023).
[13] Deci, "YOLO-NAS by Deci Achieves State-of-the-Art Performance on Object Detection Using Neural Architecture Search", Available online: https://deci.ai/blog/yolo-nas-object-detection-foundation-model/ (accessed on 1 March 2024).
[14] Highways Transport and Planning, "Highway Inspection Manual", 2022.