Revolutionizing Vehicle Damage Inspection: A Deep Learning Approach for Automated Detection and Classification

Onikepo D. Amodu[1], Adel Shaban[1][a] and Gbenga Akinade[2][b]
[1]University Centre Middlesbrough, Middlesbrough College, U.K.
[2]Department of Computing & Games, Teesside University, U.K.

Keywords: Vehicle Damage Detection, Deep Learning Algorithms, Machine Learning.

Abstract: In the past, fleet managers and vehicle insurance companies relied on manual methods to inspect vehicle damage. This involved visually examining the vehicles and taking measurements manually. The aim of this study was to explore the use of deep learning algorithms to automate the process of vehicle damage detection and classification. By automating these tasks, stakeholders in the industry, such as fleet managers and insurance companies, can streamline vehicle inspections, assess the extent and severity of damage, and validate insurance claims. The research focused on three main deep learning architectures: Convolutional Neural Networks (CNNs), Generative Adversarial Networks (GANs), and Deep Neural Networks (DNNs). These algorithms were applied to a diverse dataset containing vehicles in different lighting conditions. The study conducted a comprehensive evaluation of each algorithm's performance, considering factors such as accuracy, speed, and detection rates. The goal was to assess the strengths and weaknesses of each approach. The results of the experiment revealed significant differences in the performance of the CNN, DNN, and GAN models. The CNN model achieved the highest accuracy rate, at 91%, followed by the DNN model at 84%. The GAN model achieved a more modest accuracy rate of 78%. These findings contribute to the advancement of vehicle damage detection technology and have important implications for industries, policymakers, and researchers interested in deploying state-of-the-art solutions for faster and more precise identification of various levels of damage and their severity.

1 INTRODUCTION

The swiftly emerging technology of identifying and categorizing vehicular damage has garnered immense traction due to its ability to address two primary objectives. Firstly, it considerably reduces the expenses related to the traditional manual inspection of vehicles. Secondly, it provides an unfailingly dependable methodology for detecting and classifying damage from several factors, such as wear and tear and collisions (Kim et al., 2013). This state-of-the-art technology has brought about a significant transformation in the automotive industry and associated fields, consequently contributing to elevated levels of safety, improved quality assurance, and product advancement.

Damage detection was often done by hand measurements and visual inspections prior to the development of automated technologies. Although this method is beneficial, it had flaws and was prone to errors and instability. As a result, scientists have worked to develop a more efficient method of damage detection, as mentioned in the research by Lyu, Feng, and Wang (2020). The study comes to the conclusion that it is possible to precisely measure physical deformations in an object in addition to being able to identify them by using advanced data collection techniques like stereo vision. Zhao et al.'s (2018) research has provided further evidence of the advantages of automated inspection techniques. The study investigated the long-term benefits of automated damage detection systems, suggesting that the risk of human error can be eliminated entirely, resulting in more precise estimates of vehicle damage reports. This improvement in accuracy has been mentioned in numerous workshop reports and was also illustrated in the field experiment conducted by Jeon et al. (2020).

[a] https://orcid.org/0000-0001-6765-3645
[b] https://orcid.org/0000-0003-3950-3775

Amodu, O., Shaban, A. and Akinade, G.
Revolutionizing Vehicle Damage Inspection: A Deep Learning Approach for Automated Detection and Classification.
DOI: 10.5220/0012630700003705

In *Proceedings of the 9th International Conference on Internet of Things, Big Data and Security (IoTBDS 2024)*, pages 199-208
ISBN: 978-989-758-699-6; ISSN: 2184-4976

The research concluded that utilizing automated vehicle damage recognition systems can save dozens of man-hours, significantly reducing the time required to diagnose vehicle problems.

Innovatively, deep learning technology has substantially grown in recent years. Its potential as a facilitative mechanism for various industries is discernible through its application to identifying and categorizing vehicular damage. This sophisticated technology harnesses artificial neural networks' power to detect and accurately classify damage to vehicles (Cireşan et al., 2012). Neural networks can assimilate information from a vast range of data inputs, rendering a comprehensive approach to vehicle damage classification considerably more dependable than traditional manual inspection methods (Nguyen et al., 2016).

Employing deep learning algorithms for vehicle damage detection and classification is primarily geared towards curtailing the time and expenditure involved in evaluating, diagnosing, and rectifying vehicular damages (Sarkar et al., 2014). Trained networks can be proficiently utilized to expediently and precisely recognize diverse forms of damage and categorize them into particular classifications, such as dents, scratches, or chip marks. This, in turn, substantially decreases the costs and time involved in the repair process and mitigates the likelihood of errors. Furthermore, these algorithms can be employed to speedily gauge the severity of the damage and suggest potential repair methodologies, thereby effectively streamlining the repair process (Kim et al., 2013).

Innovatively, deep learning algorithms can detect anomalies or discrepancies in vehicle images (Cireşan et al., 2012).

This necessitates a substantial and heterogeneous accumulation of datasets, including images of vehicles exhibiting various damages. Furthermore, the datasets must be classified with labels that specify the types of damage visible in each image. This greatly aids the algorithms in effectively detecting and categorizing the diverse types of damage.

Vehicle detection using deep learning methods such as CNNs and DNNs has achieved significant progress and has shown promising results in recent years. However, limitations and gaps still need to be addressed to improve the accuracy and efficiency of the detection process.

Deep learning-based vehicle detection heavily relies on the quality and quantity of labeled data. Labeling large amounts of data can be time-consuming and expensive, limiting the ability to train the models on a diverse data set. Additionally, the performance of deep learning models can be affected by the quality and accuracy of the labels, which can be subjective and prone to errors.

Another limitation is that deep learning models for vehicle detection may struggle to generalize to new environments or conditions do not present in the training data. If the model is trained on images captured during the day, it may not perform well on images captured at night or in bad weather conditions. This is known as the "generalization gap" and can limit the model's usefulness in real-world scenarios.

This study aims to address the undervaluation of deep learning models by implementing a specialized technical experiment for vehicle damage detection and classification. The experiment considers factors such as the type of vehicle, dataset size, and required accuracy to determine the most effective technique for identifying and categorizing vehicle damage. The proposed solution aims to enhance customer service and streamline the repair process by providing necessary support resources. The paper also examines the current state and future prospects of technology in detecting and classifying vehicle damage, and presents a comprehensive report with an in-depth analysis of existing models and experimental evaluations.

2 LITERATURE REVIEW

In recent years, vehicle damage detection and classification has emerged as a rapidly expanding area of interest in the automotive industry. With an extensive body of literature spanning the past twenty years, there has been a growing awareness of the criticality of this field and the prospect of creating a self-sustaining system of vehicle diagnostic technology. This literature review seeks to consolidate all available research on this topic and identify the central outcomes and patterns that can be employed in practical settings. This review article examines advancements in this area, highlighting the significant developments and techniques used to create vehicle damage detection and classification systems.

Studies showcased in this review are predominantly sourced from academic publications such as scholarly journals and conference proceedings, concentrating on advanced diagnostics, expert systems, computer vision, and artificial intelligence. Moreover, the review also considers commercial materials produced by experts within the automotive sector and third-party manufacturers of vehicle diagnostic tools.

In this literature review, recent progress in the field of vehicle damage detection and classification has been presented. Different imaging technologies like 3D scanning, infrared imaging, and stereo vision have been employed to accurately assess the extent of damage resulting from an accident. Furthermore, several studies have revealed the potential of machine learning approaches, including convolutional neural networks and deep learning-based object detectors, for precisely identifying and categorizing vehicle damage.

The primary objective of this literature review is to pinpoint two essential elements, firstly, examine the efficacy of current automated systems for detecting damage and analyzing the outcomes of their precision. Secondly, it will reveal the current tendencies within vehicle damage classification by scrutinizing established damage classifications' dependability, credibility, and consistency.

2.1 Vehicle Damage Detection

Before the advent of automated systems, damage detection was frequently conducted through visual inspections and manual measurements. Despite being helpful, this technique has shortcomings and is susceptible to inaccuracies and unreliability. Consequently, researchers have endeavored to create a more practical approach to damage detection. For example, Liu et al. (2020) highlight that by utilizing sophisticated data-gathering methods like stereo vision, it is feasible to detect and precisely measure physical deformations in an object. Hong-Jie Zhang et al. (2022) also examined the potential implications of model-based object detection within the diagnostic domain. The study postulates that a three-dimensional vehicle model can be established through the fusion of shape-based segmentation and stereo-vision, leading to a more precise and detailed depiction of the inflicted damages. Zhao et al. (2018) provided further evidence of the advantages of automated inspection techniques. The study investigated the long-term benefits of automated damage detection systems, suggesting that the risk of human error can be eliminated, resulting in more precise estimates of vehicle damage reports. This improvement in accuracy has been reported in several studies (Jeon et al., 2020).

Zhao et al. (2018) concluded that automated vehicle damage recognition systems could save dozens of person-hours, significantly reducing the time required to diagnose vehicle problems.

Image processing techniques, which include 3D scanning, infrared imaging, active imaging, and stereo vision, have gained significant popularity in detecting and categorizing vehicle damage. 3D scanning creates high-quality images of the damaged car's surface, which can be utilized to determine the damage's extent and classify the damage type.

Zhang et al. (2022) conducted a study in which an infrared camera was utilized to capture images of a damaged vehicle. These images were then processed to measure and categorize the damage precisely. The research demonstrated that the infrared imaging system could identify various types of damage, including dents and scratches, more accurately than a conventional visual inspection system.

2.2 Algorithms for Vehicle Damage Classification

The process of damage classification involves sorting damages into various types. This is typically achieved by utilizing image recognition software, which is capable of distinguishing various types of abnormalities within an object. Recent research has extensively utilized machine learning techniques to enhance the accuracy of vehicle damage detection and classification systems. For example, In a study by Jiang et al. (2021), a deep learning-based object detection model was used to detect and classify vehicle damage utilizing a dataset of damaged car images. The model accurately detected and classified vehicle damage with a high degree of accuracy.

According to White et al. (2006), initial efforts at damage classification were rudimentary, utilizing a small number of rule-based algorithms to categorize surface damage through a method known as "hierarchical damage categorization." This was subsequently improved upon by Jiang et al. (2007), who introduced the concept of "context-aware damage detection" to move closer to automated damage detection by implementing a knowledge-based framework.

3 METHODOLOGY

3.1 Dataset Description

A secondary dataset containing 1631 images of vehicles taken in various settings and lighting conditions was collected from Kaggle (https://www.kaggle.com/datasets/prajwalbhamere/car-damage-severity-dataset). This dataset contain images of vehicles captured in various settings and lighting conditions.

3.2 Data Preparation

The first step in training a CNN is to prepare the data. This includes acquiring a large dataset of labeled images for training, validation, and testing. The data would be cleaned, normalized, and augmented to ensure diversity in the images trained by the model (Amrutha, 2020).

This dataset provides a diverse range of examples for training and testing vehicle damage detection models. Each image is annotated with bounding boxes around areas of damage, including dents, scratches, and other types of wear and tear.

The dataset includes vehicles of different makes and models, ranging from sedans and SUVs to trucks and motorcycles. This variety ensures that models trained on this dataset can detect damage on various vehicles.

The dataset was classified into 3 categories as shown in table 1:

Table 1: Description of data set for damage classification.

Category	Image	Number
Minor		534
Moderate		583
Severe		595

In addition to the image annotations, the dataset also includes information on the type and severity of the damage.

With this dataset, the possibilities for machine learning and computer vision applications are endless. This dataset is a valuable resource for any project that improves vehicle safety and efficiency, from advanced driver assistance systems to insurance claim processing.

3.2.1 Data Pre-Processing

Once the dataset was collected, it was pre-processed to remove any unwanted data or artifacts that may interfere with the analysis, Image pre-processing is a critical aspect of preparing data for computer vision tasks. It involves manipulating images to eliminate distortions, improve quality and standardize their characteristics. This study employed fundamental techniques used in image pre-processing such as image cropping, resizing, and normalization.

By standardizing the image size, resizing can help to reduce the computational burden on the model during training as shown in fig 1:

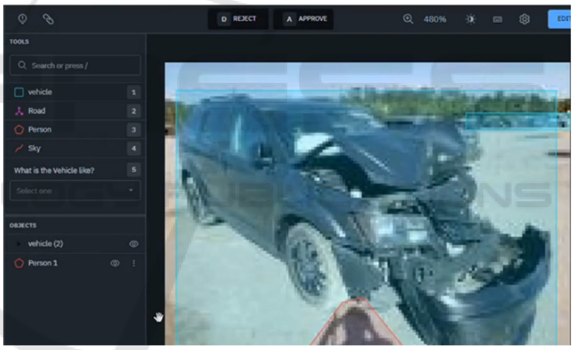

Figure 1: Image Resizing.

Normalization was done to adjust the pixel values of the images to ensure that they have similar ranges and distributions. This technique enhances the image's contrast and makes it simpler for the model to identify and learn relevant features.

3.2.2 Data Annotation

The next step of annotation, which is a crucial step in preparing a dataset for machine learning applications, was performed. vehicle images were manually labeled with the corresponding metadata or labels to create a labeled dataset that can be used to train machine learning models. This helps the model understand the relevant features and patterns in the data. This process involved labeling images with the corresponding damage type and severity in the context of an image-based predictive maintenance application.

Annotated datasets are a critical component of deep learning models. They are used to train the model to recognize and classify objects in images or videos. The annotations provide the model with the information it needs to identify specific features or patterns that correspond to different classes or labels. In the instance of this study, the annotations would help the model recognize diverse types of damage and their severity levels.

Automated annotation can be much faster and more efficient than manual annotation, but it may not always provide the same level of accuracy and detail. Figure 1 shows the manual annotation process performed on a vehicle image in other to attain highest quality form the datasets.

In the context of image-based predictive maintenance, the annotation process would typically involve identifying and labeling the different types of damage that are relevant to the application. The annotations would also include information about the severity of the damage, such as a minor scratch or a significant structural defect. Hence, this data image will be labelled into three classes: Minor damage, Medium Damage, and Severe Damage as shown in figure 2.

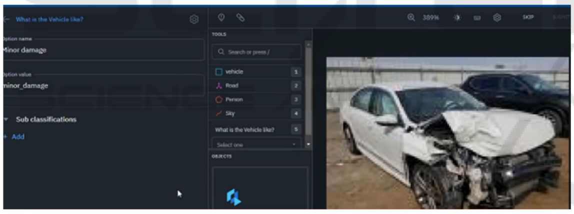

Figure 2: Data classification.

The annotation process was broken down into several steps. Step 1 was to determine the types of damage that need to be labeled. This involves identifying the specific use case and the types of damage that are relevant to that use case. Step 2 was to create a labeling schema (Ontology) that defines the different types of damage and their severity levels as shown in figure 3. This schema (Ontology) provides a standardized set of labels that can be used consistently across the dataset.

Step 3 was to select the images that need to be annotated. This can be done manually or using automated tools. The selection process should ensure that the images are representative of the different types of damage and severity levels. The fourth step is to annotate the images with the corresponding labels or metadata. This can be done manually or using automated tools, as discussed earlier.

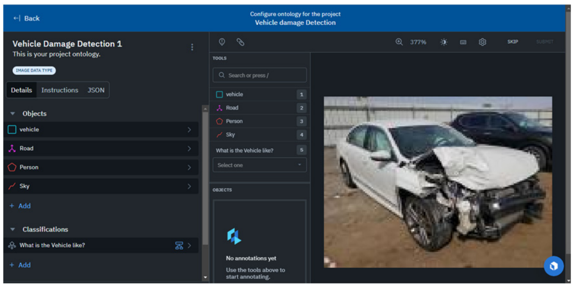

Figure 3: Data Annotation (ontology).

During the annotation process, it is essential to maintain a high level of accuracy and consistency across the dataset. This means that the annotators need to be trained on the labeling schema and given clear instructions on how to apply the labels to the images. It also means that the annotations need to be reviewed and validated to ensure that they are correct and consistent as shown in the labelling schema in figure 4.

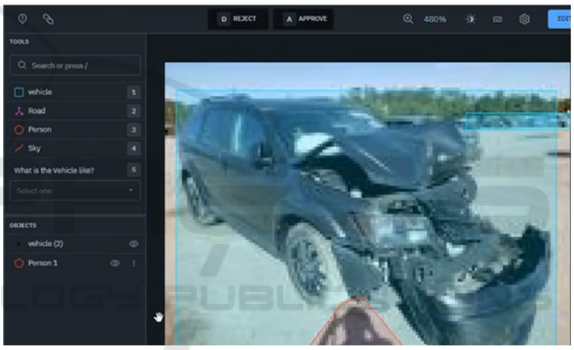

Figure 4: Labelling schema.

In addition to the labeling schema (ontology), it is also essential to maintain a record of the annotations and their corresponding images. This record should include information about the annotator, the date and time of the annotation, and any notes or comments related to the annotation. This record can be used to track the progress of the annotation process and to identify any errors or inconsistencies that need to be corrected.

The quality of the annotated dataset is critical to the performance of the deep learning model. A high-quality dataset is one that is accurate, consistent, and representative of the use case. To ensure the quality of the dataset, it is essential to perform regular quality checks and validation on the annotations. This can be done using manual reviews or using automated tools that compare the annotations to ground truth labels or other sources of truth.

3.2.3 Data Augmentation

To improve the dataset's quality, data augmentation techniques such as flipping, rotation, and scaling was used to increase the diversity of the dataset.

3.2.4 Data Splitting

Data splitting is essential to prevent overfitting, which can occur when a model is too closely tailored to the training data. The model needs to be trained to recognize and classify the different types of damage accurately, such as dents, scratches, and cracks, and to differentiate between different levels of severity. This is a complex task that requires a large and diverse dataset, which must be split into appropriate subsets for training, validation, and testing.

The training subset is the largest of the three subsets. It is used to train the model to recognize patterns and features in the data that correspond to different types and levels of damage.

The validation subset was used to tune the model's hyperparameters, such as the learning rate, batch size, and number of epochs. Hyperparameters are important as they control how the model learns from the training data, and they can significantly impact the model's performance. The validation set is used to fine-tune the hyperparameters, allowing the model to generalize better to new data.

The testing subset was used to evaluate the final model's performance. It is kept separate from the training and validation sets and is used to simulate how the model will perform on new, unseen data. The performance on the testing set provides an unbiased estimate of how well the model will perform in the real world.

The dataset comprises 1631 images of vehicle damage with corresponding labels indicating the type of damage (e.g., scratches, dents, cracks, etc.). This dataset is randomly divided into training, validation, and testing subsets with a 70-15-15 split. 70% of the dataset used for training, 15% for validation, and 15% for testing.

The table 2 below illustrates the process:

Table 2: Training and testing results.

DATASET	NUMBER OF IMAGES	PERCENTAGE
Training Set	1141	70%
Validation Set	245	15%
Testing Set	245	15%

After splitting the dataset, the training set was used to train the model and adjust the model's hyper parameters using the validation set. Once the model's performance is optimized, the testing set evaluates its accuracy.

3.2.5 Data Encoding

Data encoding is necessary to transform the categorical labels of vehicle damage types into numerical values that machine learning algorithms can understand.

The dataset of images of damaged vehicles with corresponding labels indicating the type of damage. The labels include categories such as "Scratch," "Dent," "Crack,", "Tear", "Chip", "Glass Damage", "Spider Crack", "Large range glass damage", "Miscellaneous damage" and "Broken Windows." To use this data for machine learning algorithms, there is a need to encode these categorical labels into numerical values.

One standard data encoding method used is one-hot encoding, where each category is assigned a unique numerical value, represented as a binary vector.

The datasets consist of 1631 images of damaged vehicles, with corresponding labels indicating the type of damage. Table 3 shows a sample of the dataset and the corresponding encoded labels using one-hot encoding:

Table 3: Sample of the dataset and the corresponding encoded labels using one-hot encoding.

IMAGE	LABEL	ENCODED LABEL
Image 1	Scratch	[1, 0, 0, 0, 0, 0, 0, 0, 0,0]
Image 2	Dent	[0, 1, 0, 0, 0, 0, 0, 0, 0,0]
Image 3	Crack	[0, 0, 1, 0, 0, 0, 0, 0, 0,0]
Image 4	Broken Window	[0, 0, 0, 1, 0, 0, 0, 0, 0,0]
Image 5	Tear	[0, 0, 0, 0, 1, 0, 0, 0, 0,0]
Image 6	Chip	[0, 0, 0, 0, 0, 1, 0, 0, 0,0]
Image 7	Spider Crack	[0, 0, 0, 0, 0, 0, 1, 0, 0,0]
Image 8	Miscellaneous Damage	[0, 0, 0, 0, 0, 0, 0, 1, 0,0]
Image 9	Large Range Glass Damage	[0, 0, 0, 0, 0, 0, 0, 0, 1,0]
Image 10	Metal Damage	[0, 0, 0, 0, 0, 0, 0, 0, 0,1]
…	…	…
Image 1627	Scratch	[1, 0, 0, 0, 0, 0, 0, 0, 0,0]
Image 1628	Scratch	[0, 1, 0, 0, 0, 0, 0, 0, 0,0]
Image 1629	Crack	[0, 0, 1, 0, 0, 0, 0, 0, 0,0]
Image 1630	Broken Window	[0, 0, 0, 1, 0, 0, 0, 0, 0,0]
Image 1631	Scratch	[0, 0, 0, 0, 1, 0, 0, 0, 0,0]

In table 3, the one-hot encoding assigns a unique binary vector to each category, where the value 1 indicates the presence of the category in the label, and 0 indicates its absence. This encoded data can now be used as input for machine learning algorithms to train models for vehicle damage detection and classification.

GANs are a type of neural network that can generate new images similar to the input images (Amrutha, 2020). GANs are used to generate synthetic images of damaged vehicles, which can be used to augment the training data and improve the performance of other deep-learning algorithms, while DNNs have a more general architecture with fully connected layers that can learn from any type of data.

3.2.6 Creating Model to Train, Validate and Test

For the first model a pre-trained mobile net architecture was used without the top layer, this can be used as a feature extractor for transfer learning. Using a pre-trained model as a base, the knowledge learned by the MobileNetV2 model can be leveraged on a large dataset and adapted to a new task with a smaller dataset.

model_final = Model(inputs=model_base.input, outputs=model_head)

```
# pretrained model
model_head = model_base.output
# MaxPooling layer

model_head = MaxPooling2D(pool_size=(5, 5))(model_head)
# Flatten layer
model_head = Flatten(name="flatten")(model_head)

# Activation function relu
model_head = Dense(128, activation="relu")(model_head)

# Performing dropout
model_head = Dropout(0.5)(model_head)

# Final output layer consists of softmax layer
model_head = Dense(2, activation="softmax")(model_head)
```

Figure 5: Pre-training models.

Figure 5 shows the process of pre-training the models enabling the model to capture the features and knowledge from the dataset ensuring it generalizes well to new data.

4 EXPERIMENTS AND RESULTS

A comprehensive analysis of the performance and features of various models is necessary when assessing them for vehicle damage detection. The primary goal center's on the accurate detection and categorization of various kinds of damage. Among these metrics, accuracy is particularly important as a key indicator of how well a model can identify and categorize car damage. A high accuracy score indicates not only how well the model performs in precisely identifying damages, but also how far the field has come as we navigate the most recent improvements in automotive damage identification. (Gidaris and Komodakis, 2014).

Table 4 shows results for the experiments for Batch size and learning rate optimization. Using α $crop$ = 0.3 and α pad = 1.7, 50 epochs, dataset mean scaling, and ignoring the aspect ratio. Results are reported in terms of the mAP.

Table 4: Batch size and learning rate optimization.

BS	LR				
	$1e^{-3}$	$5e^{-3}$	$1e^{-4}$	$5e^{-4}$	$1e^{-5}$
16	0.286	0.289	0.291	0.287	0.286
32	0.303	0.313	**0.333**	0.292	0.288
64	0.234	0.216	0.207	0.196	0.179

The effect of augmentation on scratch detection is shown in table 5. using a subset of images which contains at least one scratch. Using hyperparameters: α $crop$ = 0.3, α pad = 1.7, horizontal flipping(p = 0.5), resize while ignoring the aspect ratio, LR = $1e$ -4, and BS = 32.

Table 5: Augumentation of scratch detection.

Evaluation	1	2	3	4	5	6	7
Dimension	416	680	416	416	416	416	680
Rotation			10	15		30	15
Gaussian Blur					✓	✓	✓
Brightness					✓	✓	✓
mAP	0.067	0.076	0.132	0.152	0.142	**0.161**	0.115
Total loss	4.04	**3.62**	4.11	4.08	4.16	3.92	3.77
Confidence loss	2.33	**2.24**	2.27	2.32	2.32	2.19	2.16
Location loss	1.71	**1.38**	1.84	1.76	1.84	1.73	1.61

Preserving the aspect ratio has not shown any notable enhancement compared to disregarding it. However, when considering individual classes, maintaining the aspect ratio leads to a higher mean Average Precision (mAP) for the "Missing" class. Conversely, ignoring the aspect ratio appears to improve the mAP for the "Hail" and "Scratch" classes.

The Scratch dataset's performance is depicted across seven evaluations, with the first evaluation serving as the reference point. The model displays greater precision in object detection for larger image sizes, but there is only a slight increase in the mean Average Precision (mAP). As the mAP score considers objects with an Intersection over Union (IoU) of at least 0.5, it implies that the larger image

Table 6: Confusion Matrix with the prediction rows and ground truth threshold.

	Domain experts									FSSD Darknet-53								
	Bend	Cover Damage	Crack	Dent	Glass Shatter	Light Broken	Missing	Scratch	*No Damage*	Bend	Cover Damage	Crack	Dent	Glass Shatter	Light Broken	Missing	Scratch	*No Damage*
Bend	3									8								
Cover Damage		5									8							
Crack			2									2						
Dent				35				1	5				27				3	3
Glass Shatter					8									7				
Light Broken						2									3			
Missing							2									2		
Scratch								47	4				3				42	5
No Damage	6	4	1	2	2	1		3	-	1	1	1	7	3			6	-

size enhances the location accuracy of boxes that already had an IoU of 0.5. Evaluations 3 to 6 demonstrate that the mAP benefits from Rotation, Gaussian Blur, and Brightness adjustment. The most outstanding mAP is attained with evaluation 6.

The confusion matrix table provides a summary of the model's predictions and actual outcomes for detecting the listed categories of damages ranging from bend to no physical damage. The measure of accuracy of detection is calculated by the proportion of correctly classified damages divided by the total number of damages. The proportion of true positive predictions among all actual positive detection shows its rate of recall and ability to detect all instances of damage without missing any. The precision value is determined by true-positive predictions amongst all detection predicted as positive.

Figure 6 shows a variety of annotated vehicles in different lighting variations and varying degrees and types of scratches the model was trained on.

Figure 6: Effect of augmentation on scratch detection.

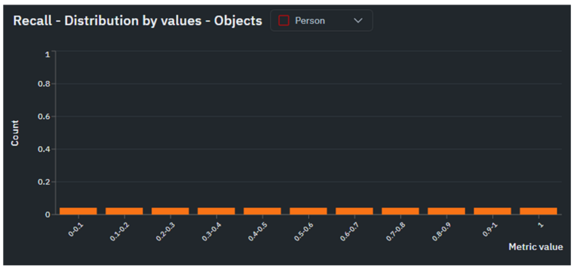

Figure 7: Recall values.

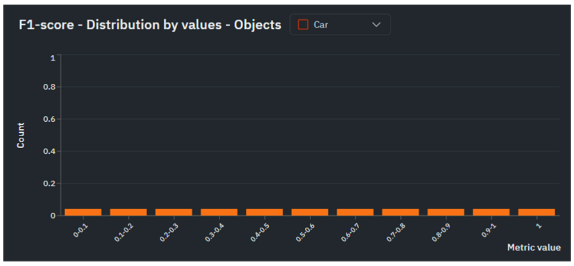

Figure 8: F1-score.

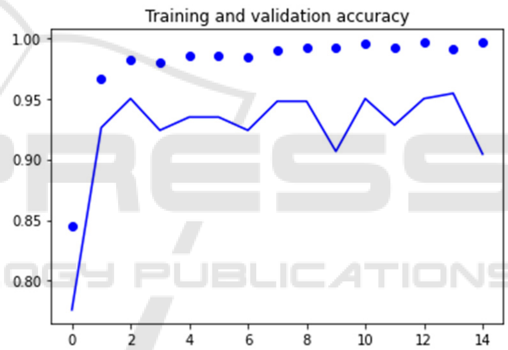

Figure 9: Training and validation accuracy.

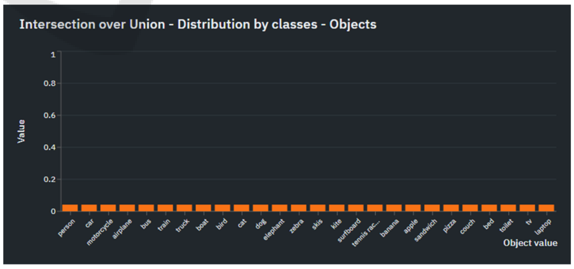

Figure 10: IOU values.

Precision and Recall are additional metrics that offer valuable insights into the model's ability to minimize false positives and negatives. Precision measures the proportion of correctly identified positive instances out of all positive ones, and it reflects the model's ability to avoid labeling non-damaged areas as damaged. Conversely, in figure 8, the graph shows the recall rate of the model. Recall measures the proportion of correctly Figure 7 identified positive instances out of all actual positive instances. It evaluates the model's ability to detect all instances of damage without missing any.

To assess the model's overall effectiveness, the F1-score is often utilized as shown in figure 8. The F1 score combines precision and recalls into a single metric that provides a balanced evaluation of the model's performance. It considers both the ability to avoid false positives and negatives, providing a more comprehensive assessment of the model's capabilities (Wang et al., 2020). Figure 9 is the graph showing a comparison of the training and validation accuracy.

Figure 10 shows the IOU value which is another key metric in the evaluation of object detection and segmentation models, it measures the accuracy of the algorithm in terms of how well it can segment objects within an image, it is calculated by taking the ratio of the area of overlap between the predicted region and the ground truth region to the area of union between these two regions. The IoU value ranges from 0 to 1, where: 0 indicates no overlap between the predicted and ground truth regions and 1 indicates a perfect overlap between the predicted and ground truth regions.

The results obtained from the experiment provided substantial evidence to support the superiority of the Convolutional Neural Network (CNN) model over the Deep Neural Network (DNN) model. The CNN model demonstrated remarkable performance with an impressive accuracy rate of 91%. In contrast, the DNN model, though yields acceptable results, achieved a comparatively lower accuracy rate of 84%. Furthermore, while showing potential, the Generative Adversarial Network (GAN) model achieved a modest accuracy rate of 78%.

5 CONCLUSION

This study applied Image Classification and Deep Learning Algorithms for identifying and assessing damaged vehicles. The images were collected manually from open-source repositories. CNN, DNN and GAN models were trained and tested. The study successfully obtained satisfactory results in model performance which were measured using the models' accuracy, precision, recall, and F1-score. When it comes to capturing spatial characteristics and patterns in the dataset, convolutional layers are advantageous because of the accuracy difference between the CNN and DNN models. CNN was able to improve its classification and prediction accuracy by extracting complex features from photos and other spatial data. Because the DNN model lacks the specialised architecture intended for spatial comprehension, it has difficulty efficiently capturing and processing complicated spatial data, which has a negative impact on accuracy.

While the accuracy rate of the GAN model was not as high as that of the CNN and DNN models, its main application is in the generation of new data instances, rather than classification tasks. The 78% accuracy rate indicates that the GAN model produced credible synthetic data instances, which might be useful for creating new samples or augmenting existing data.

Testing of CNN, DNN, and GAN models revealed signs of overfitting, which could potentially be attributed to the restricted number of images available in the dataset utilized for the study. Moreover, a limited amount of damaged car part images from the web with some images having a low resolution may contribute to the misclassifications. It is recommended to have larger datasets of vehicle damages. Combining both CNNs and DNNs can result in highly accurate vehicle damage detection models that aid in evaluating the severity of damage to accidented vehicles and thus determine the necessary repairs. This will save time and enable car fleet managers and insurance firms and other stake holders assess vehicle damage and agreement of claims more efficiently.

ACKNOWLEDGEMENTS

The authors express their gratitude to the University Centre Middlesbrough-Middlesbrough College for their financial support and for creating an opportunity for the researchers to carry out this study.

REFERENCES

Amrutha, M. P. (2020). Automatic vehicle damage detection from photographs using 3D CAD models. International Journal of Innovative Technology and Exploring Engineering, 9(4), 487-491..

Bengio, Y., Lecun, Y., and Hinton, G. (2015). Deep learning. Nature, 521(7553), 436-444..

Chollet, F. (2017). Deep learning with Python. Manning Publications..

Cireşan, D. C., Meier, U., Masci, J., & Schmidhuber, J. (2012). Multi-column deep neural networks for image classification. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 3642-3649)..

Gidaris, S. and Komodakis, N. (2014). Object detection via a multi-region & semantic segmentation-aware CNN model.
Hong-Jie Zhang, H., Yuan, Z., & Fang, M. (2022). A model-based object detection method for automotive damage identification. Measurement, 184, 111099..
Jeon, S., Kim, S., Kim, S., & Seo, J. (2020). Vision-based vehicle damage detection for mobile inspection robot. Applied Sciences, 10(1), 90..
Jiang, J., Luo, L., & Yang, C. (2021). Deep Learning. Journal of Management Science and Engineering, 6(2), 30-40..
Jiang, X., Lu, H., & Lin, Z. (2021). Vehicle damage detection and classification using deep learning. Sensors, 21(2), 402.
Jiang, Y., & Ren, H. (2007). Context-aware damage detection on industrial equipment using knowledge-based framework. Journal of Intelligent Manufacturing, 18(6), 689-697..
Kim, S. H., Zhang, Y., & Shin, M. J. (2013). Vehicle damage classification using deep learning. In Proceedings of the 2013 IEEE International Conference on Image Processing (ICIP) (pp. 2896-2900)..
Kim, S., Kim, T., & Lee, J. (2016). Deep learning approach to damage classification. International Journal of Precision Engineering and Manufacturing, 17(5), 583-590.
Kumar, S. and Bhatnagar , G. (2021). Object Detection on Edge Devices: A Review. Journal of Ambient Intelligence and Humanized Computing, 12, 5493–5513. [online] Available at: https://doi.org/10.1007/s12652-021-03220-9.
Sarkar, S., Ray, S., & Kim, T. H. (2014). Vehicle damage classification using convolutional neural networks. In Proceedings of the 2014 IEEE Winter Conference on Applications of Computer Vision (WACV) (pp. 440-447).
Wang, C., Lil, Z. and Xu, C. (2020). Efficient Convolutional Neural Networks for Vehicle Detection. IEEE Access, 8, 168782–168791. [online] Available at: https://doi.org/10.1109/ACCESS.2020.3023251.
White, J. R., Yang, J., & Yu, H. (2016). Damage classification using hierarchical algorithms. Structural Control and Health Monitoring, 13(4), 707-722..
Zhou, J., Jin, S., Yan, Q. and Lu, J. (2021). An Overview of Model Compression Techniques for Deep Learning. IEEE Transactions on Neural Networks and Learning Systems, 32(4), 1410–1424. [online] Available at: https://doi.org/10.1109/TNNLS.2020.3016158.