

ENHANCED VEHICLE DAMAGE CLASSIFICATION USING ATTENTION-AUGMENTED DEEP LEARNING ON THREE-QUARTERS VIEW CAR DAMAGE DATASET

Thesis Submitted
in Partial Fulfillment of the Requirements for the
Degree of

MASTER OF TECHNOLOGY
in
Software Engineering

by
Krithic S
(24/SWE/20)

Under the Supervision of
Dr. Divyashikha Sethia
Associate Professor, Department of Software Engineering
Delhi Technological University



To the
Department of Software Engineering
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi-110042, India

May, 2026

ACKNOWLEDGEMENT

In the first place, it is necessary for me to offer my sincere appreciation to my supervisor and project guide **Dr. Divyashikha Sethia** for the valuable guidance, support, and encouragement she gave me during the entire period of this project.

I am also deeply grateful to everyone else who contributed to the successful completion of this report through their support and assistance.

(**Krithic S**)

24/SWE/20

Department of Software Engineering
Delhi Technological University



DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi-42

CANDIDATE'S DECLARATION

I **Krithic S** hereby certify that the work which is being presented in the thesis entitled **Enhanced Vehicle Damage Classification Using Attention-Augmented Deep Learning on Three Quarters View Car Damage Dataset** in partial fulfillment of the requirements for the award of the Masters Degree, submitted in the Department of **Software Engineering**, Delhi Technological University is an authentic record of my own work carried out during the period from **2025** to **2026** under the supervision of **Dr Divyashikha Sethia**.

The matter presented in the thesis has not been submitted by me for the award of any other degree of this or any other Institute.

A handwritten signature in blue ink, appearing to be 'Krithic S', written over a horizontal line.

Candidate's Signature

This is to certify that the student has incorporated all the corrections suggested by the examiners in the thesis and the statement made by the candidate is correct to the best of our knowledge.

A handwritten signature in blue ink, appearing to be 'D. Mir', written over a horizontal line.

Signature of Supervisor (s)

Signature of External Examiner



DELHI TECHNOLOGICAL UNIVERSITY

(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi-42

CERTIFICATE BY THE SUPERVISOR

I hereby certify that Krithic S (Roll No. 24/SWE/20) has carried out their research work presented in this thesis entitled “Enhanced Vehicle Damage Classification Using Attention-Augmented Deep Learning on Three Quarters View Car Damage Dataset” for the award of Master of Technology from The Department of Software Engineering, Delhi Technological University, Delhi under my supervision. The thesis embodies results of original work, and studies are carried out by the student himself and the contents of the thesis do not form the basis for the award of any other degree to the candidate or to anybody else from this or any other University/Institution. To the best of my knowledge this work has not been submitted in part or full for any Degree or Diploma to this University or elsewhere.

Place: Delhi

Date: 19.05.2026

A handwritten signature in blue ink, appearing to read "DWS", with a long horizontal line extending from the end.

Dr. Divyashikha Sethia
Associate Professor
Dept. of Software Engineering
Delhi Technological University
New Delhi, India

ABSTRACT

The manual detection process for vehicle damage is monotonous and prone to errors. Especially when dealing with massive-scale scenarios like insurance claims assessment and repair costs evaluation, the process becomes cumbersome. In this report, we propose a classification framework that automatically identifies damage using image processing techniques and deep learning algorithms. We evaluate our proposed system using the TQVCD dataset that includes images from a three-quarter view showing both frontal and lateral parts of vehicle damage.

All images go through a pre-processing stage where we resize, normalize, and randomly augment them by rotating, horizontal flipping, and adjusting their brightness levels. We compare three architectures in terms of performance on TQVCD dataset: EfficientNet-B3, DenseNet169, and ResNet50 augmented with an attention mechanism. In the latter, the attention module allows the network to attend only to the relevant part of the image where vehicle damage is located. On TQVCD, ResNet50 with Attention achieves 87% accuracy, DenseNet169 achieves 85%, and EfficientNet-B3 achieves 82%.

Another dataset used to verify the effectiveness of our proposed architecture in different vehicle categories is CarDD, which includes many types of damages. On CarDD, ResNet50 with Attention achieves 85% accuracy, while DenseNet169 and EfficientNet-B3 achieve 83% and 80% respectively. Lastly, on TartesiaDS, an insurance-oriented dataset with annotations made by professionals under various lighting conditions, ResNet50 with Attention achieves 89% accuracy, which is higher than the baseline ResNet50 architecture tested on the same dataset reported by Pérez-Zarate et al. (2024) achieving 85.8%. On TartesiaDS, DenseNet169 achieves 87% and EfficientNet-B3 achieves 83

Overall, our study proves the reliability and effectiveness of the attention-augmented deep learning models in the process of automating damage detection in vehicles. In particular, on the three presented benchmarks, architectures with attention modules perform better than their counterparts without such mechanisms. Our future research will focus on ensemble fusion of multiple models, grading the degree of damage, and utilizing explainable AI methods like Grad-CAM.

Contents

1	Introduction	4
1.1	Background	5
1.2	Problem Statement	5
1.3	Objectives	6
2	Technical Background	7
2.1	Deep Learning in Vehicle Damage Classification	7
2.1.1	Convolutional Neural Networks (CNNs)	7
2.1.2	Transfer Learning	8
2.1.3	Model Architectures	8
2.1.4	Optimization and Regularization	9
2.2	Evaluation Metrics	9
2.3	Dataset Preparation and Preprocessing	10
3	Literature Review	11
3.1	Overview of Vehicle Damage Detection Systems	11
3.2	Deep Learning Methods for Automotive Damage Detection	12
3.3	Multi-view Image Recognition for Damage Detection	13
4	Recent Work on Vehicle Damage Datasets	14
4.1	Overview of Vehicle Damage Datasets	14
4.2	Comparative Analysis of Recent Studies	14
4.3	Studies on TQVCD Dataset	15
4.4	Studies on CarDD Dataset	16
4.5	Studies on TartesiaDS Dataset	16
4.6	Trends in Model Architectures	17
4.7	Common Techniques and Approaches	17
4.8	Performance Analysis and Benchmarks	17
4.9	Research Gaps and Opportunities	18
4.10	Summary	19
5	Methodology	20
5.1	Dataset Description	20
5.2	ResNet50 with Attention Mechanism	21
5.3	EfficientNet-B3	23
5.4	DenseNet169	23

6	Results	25
6.1	Model Performance Comparison	25
6.2	Comparison with Previous Works	26
6.3	Sample Classification Results	27
6.4	Analysis	28
	6.4.1 Key Observations	28
6.5	Discussion of Results	29
6.6	Limitations	30
7	Conclusion and Future Work	31

List of Tables

4.1	Comparison of Recent Work on Vehicle Damage Detection Datasets .	15
6.1	Model Performance Comparison on Different Datasets	25
6.2	Comparison of Current Study with Previous Works on the Same Datasets	26

List of Figures

5.1	Architecture of ResNet50 with Attention Mechanism for vehicle damage classification.	22
5.2	Working of the DL model	24
6.1	Sample classification results of ResNet50 with Attention on the TQVCD dataset. ✓ indicates correct prediction and × indicates misclassification.	27

LIST OF ABBREVIATIONS

AI	Artificial Intelligence
AUC	Area Under the Curve
CarDD	Car Damage Detection Dataset
CNN	Convolutional Neural Network
DL	Deep Learning
FN	Front Normal
FB	Front Breakage
FC	Front Crushed
FP	False Positive
FN	False Negative
Grad-CAM	Gradient-weighted Class Activation Mapping
GPU	Graphics Processing Unit
ImageNet	Large Scale Visual Recognition Dataset
ML	Machine Learning
RB	Rear Breakage
RC	Rear Crushed
RN	Rear Normal
ReLU	Rectified Linear Unit
ResNet50	Residual Network with 50 Layers
SHAP	SHapley Additive exPlanations
TN	True Negative
TP	True Positive
TQVCD	Three-Quarter View Car Damage Dataset
TartesiaDS	Tartesia Damage Dataset
ViT	Vision Transformer
XAI	Explainable Artificial Intelligence

Chapter 1

Introduction

Classification of vehicle damages is one of the significant tasks in the industries such as automotive insurance, repairs, and accidents analysis [1]. Manual vehicle damage examination can take a long time, produce inconsistent results, and becomes hard to perform if the number of cars is big. Automating this task by means of deep learning models helps to achieve better accuracy and speed up the process, thus making the classification practical in a real-world environment [2].

CNN models and transfer learning approaches in deep learning have proved to work well in image classification tasks since these models are able to automatically learn and capture meaningful spatial features from images [1]. In this paper, the Three-Quarter View Car Damage Dataset (TQVCD) is used as the major dataset [3] to be studied since the images of this dataset represent vehicles' three-quarter view giving the broad view of possible damage [2]. In addition to TQVCD, other benchmarks, namely CarDD [4] and TartesiaDS [5] will be considered to make a comparative analysis of the proposed models using different images and damage cases.

The primary goal of the study will be to test and validate the ability of state-of-the-art deep learning algorithms, such as EfficientNet-B3 [6], DenseNet169 [7], and ResNet50 with Attention [8] in classifying damages, like scratches, dents, bumps, etc. We choose these models for their strong feature extraction capability and suitable for transfer learning-based image classification [9]. Particularly, the attention mechanism introduced into the ResNet50 network helps to focus on the relevant part of the image containing the damage [8].

Thus, in this research, it will be necessary to conduct the analysis of the selected deep learning models on multiple datasets using the basic set of evaluation metrics, namely accuracy, precision, and recall to create an efficient and scalable framework for automated vehicle damage classification [1].

1.1 Background

There is therefore the growing need to develop automated systems for vehicle damage classification to boost efficiency.

Deep learning methods can be employed successfully in image-based classification because they learn useful features from raw vehicle images automatically [1]. Such models have shown great promise in recognizing different vehicle damages such as scratches, dents, and bumps [10]. Challenges such as overfitting, lack of diversity in data, and difficulties in selecting the most important damaged regions hamper the performance of classification. [11]

Architecture innovations such as ResNet, EfficientNet, and DenseNet have addressed some of these problems [9]. ResNet facilitates deeper learning using residual connections [9], while EfficientNet enhances performance via compound scaling [6]. Additionally, DenseNet enhances feature reuse through dense connections [7]. Attention mechanisms allow the model to pay more attention to the regions where damages occur improving its performance [8].

The Three-Quarter View Car Damage Dataset (TQVCD) [3] has been adopted as the main dataset used in this research. This dataset comprises images acquired from three-quarters of the vehicle view making it very informative about visible damages [2]. The CarDD dataset [4] and TartesiaDS [5] datasets can be used as benchmarks to evaluate the performance of the proposed system. The application of these advanced deep learning models on these three.

1.2 Problem Statement

Vehicle damage assessment is an essential process that happens in the automotive industry, particularly during insurance claim processing and repairs [1]. However, traditional manual methods are generally slow, inconsistent, and error-prone, especially when there is a need to evaluate numerous vehicles [2]. This leads to delayed claims resolution, subjective damage evaluation and high operation costs [1].

Despite being efficient in automatic vehicle damage classification, the deep learning technique presents multiple challenges. Firstly, most of the approaches are focused on a specific database and simple architecture, hence limiting its generalisation capabilities in terms of the diversity of images and damages [11]. Another problem associated with this issue is ineffective detection of relevant damages which affects the classification accuracy [8].

The above problems will be addressed in this report through testing of deep learning architectures including ResNet50 with Attention [8], EfficientNet-B3 [6] and DenseNet169 [7] based on attention mechanism and transfer learning on three distinct databases, namely TQVCD [3], CarDD [4] and Tartesia [5].

1.3 Objectives

The following are the primary goals of this research:

- **Development of an Enhanced Deep Learning Model Framework:** This report aims to create a vehicle damage identification framework utilizing EfficientNet-B3 [6], DenseNet169 [7], and ResNet50 with Attention [8] to recognize scratches, dents, and bumps.
- **Application of Transfer Learning Techniques:** This report uses pre-trained deep learning models on ImageNet data [2] for improved feature extraction and better classification performance on limited datasets [1] [9].
- **Implementation of Spatial Attention:** This report aims to incorporate a spatial attention mechanism into the ResNet50 model to enable the model to focus on specific damaged areas within the car image to improve classification performance [8].
- **Evaluation of Models' Performance:** This report compares the classification efficiency of all three models based on metrics such as accuracy, precision, and recall to select the most efficient model for vehicle damage classification among TQVCD [3], CarDD [4], and TartesiaDS [5].
- **Improving Vehicle Damage Assessment:** This report creates an automated vehicle damage recognition model that improves the efficiency of vehicle damage assessment in the automotive industry [1].

Chapter 2

Technical Background

In this section, the theoretical background of deep learning for classification of vehicles' damages is discussed. It describes the key concepts of CNN, transfer learning, advance deep learning approaches, metrics, and dataset preparation methods that are used throughout the paper. Vehicle damage classification is an image-based task aimed at classifying types of vehicle damage such as scratches, dents, and bumps using images of damaged vehicles [2]. Because each of these types can differ in size and shape, deep learning models are particularly applicable to detect visual features in vehicle images for their accurate classification [1].

2.1 Deep Learning in Vehicle Damage Classification

Modern deep learning techniques have revolutionized image recognition processes, in particular when it comes to recognizing types of vehicle damage [1]. Convolutional Neural Networks are especially vital in this process since they allow for automated feature extraction from vehicle images and further classification of types of vehicle damage (scratches, dents, and bumps) [10]. These models perform hierarchical learning and, thus, are ideal for application to vehicle damage classification [1].

2.1.1 Convolutional Neural Networks (CNNs)

CNNs are structured to learn hierarchical representations of spatial features through multiple layers of convolution, pooling, and fully connected layers, as explained below:

- **Convolutional Layers:** Convolutional layers are used for detecting low-level features such as edges, textures, and other patterns necessary to detect vehicle damages [10].

- **Pooling Layers:** Pooling layers reduce the dimensionality of the input image and help in retaining only the most critical features from an image[1].
- **Fully Connected Layers:** Fully connected layers aggregate all the extracted features and generate output predictions about the damage type of the vehicle[1].
- **Softmax Layer:** Softmax generates probability scores for each class of damage for performing multi-class classification[9].

Furthermore, it is imperative to acknowledge that the utilization of convolutional neural networks has fundamentally transformed the landscape of computer vision applications, enabling unprecedented levels of feature extraction accuracy.

2.1.2 Transfer Learning

Data labelling can be challenging in the real world since collecting sufficient labelled data for any specific application is difficult. Transfer learning solves this problem by making use of the knowledge acquired by models trained on large datasets like ImageNet[9]. Transfer learning helps improve model performance even when there is insufficient data [1].

In this report, three ImageNet-trained models, namely **ResNet50**, **EfficientNet-B3**, and **DenseNet169**, will be fine-tuned to classify vehicle damages [9]. Pre-trained models like ResNet50, EfficientNet-B3, and DenseNet169 have learned generalized visual representations containing useful features such as edges, textures, and patterns required for visual recognition[6][7]. Fine-tuning these models makes them more efficient in solving a particular problem while benefiting from transfer learning [2].

2.1.3 Model Architectures

- **ResNet50:** ResNet50 [9] introduces residual connections to address the vanishing gradient issue in deep networks, making it possible to train extremely deep models that are efficient in processing complex images. Therefore, it is a viable option for vehicle damage recognition due to its effectiveness in discriminating between damage classes through visual differences.
- **DenseNet169:** In DenseNet169 [7], each layer of the network is connected to every other layer of the following blocks, which makes it effective in terms of reusing features and efficiently training deep neural networks. As a result, it is suitable for detecting fine-grained damage, such as scratches and dents in the vehicle[8].

- **EfficientNet-B3:** EfficientNet-B3 [6] employs a technique called compound scaling which allows simultaneous scaling of the network’s width, depth, and resolution while preserving optimal balance. Such technique helps in building computationally efficient models with outstanding performance, making it useful for detecting fine-grained vehicles damage classes[9].

2.1.4 Optimization and Regularization

For all three neural networks, the following techniques are applied:

- **Optimizer and Loss function:** Adam optimizer is used together with cross-entropy loss for multi-class classification with a learning rate of 0.001[9].
- **Learning Rate for ResNet50 with Attention:** A decreased learning rate of 0.0005 is applied to achieve convergence[8].
- **Dropout:** Dropout with a rate of 0.5 is applied in order to mitigate overfitting in EfficientNet-B3 and DenseNet169[6][7].
- **Batch Normalization:** Batch normalization in DenseNet169 is applied after each dense block [7].
- **Data Augmentation:** Random rotations, horizontal flips, and varying brightness[11].

2.2 Evaluation Metrics

We evaluate model performance using common classification metrics calculated from the confusion matrix [1]. Suppose that TP, TN, FP, and FN represent true positive, true negative, false positive, and false negative counts, respectively. The metrics are defined as:

- **Accuracy:** Accuracy is the ratio of correct classifications of the images of vehicles out of all images tested by the model.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (2.1)$$

- **Precision:** Precision is the ratio of correct predictions of damaged images among all predicted damaged images.

$$\text{Precision} = \frac{TP}{TP + FP} \quad (2.2)$$

- **Recall (Sensitivity):** Recall represents the ratio of real damaged images among those predicted to be damaged.

$$\text{Recall} = \frac{TP}{TP + FN} \quad (2.3)$$

2.3 Dataset Preparation and Preprocessing

In this report, all images of the TQVCD dataset [3] are resized to 224×224 pixels and pixel values are normalised to the range $[0, 1]$ [2]. To improve model robustness and reduce overfitting, we apply data augmentation techniques including random rotation, horizontal flipping, and brightness adjustment [11]. We encode class labels as numeric variables to support the multi-class classification pipeline [1]. These steps ensure the data is suitably prepared for training and help the model generalise better to unseen images [2].

Chapter 3

Literature Review

The chapter below provides an overview of past literature on vehicle damage detection systems. In particular, it highlights methodologies, data sources, and performance achieved with the various models that inform this area. It focuses on the utilization of deep learning techniques for the same purposes, particularly those based on attention-based models, ResNet50, EfficientNet-B3, DenseNet169, and multi-view image recognition techniques[8]. It also identifies existing challenges associated with vehicle damage detection systems which the proposed improvements attempt to address.

3.1 Overview of Vehicle Damage Detection Systems

Previously, inspectors performed damage assessments manually, taking time and lacking consistency due to inability to deal with many vehicles [2]. In early vehicle damage classification studies, researchers used simple computer vision techniques such as edge detection and template matching [11]. Unfortunately, such techniques proved incapable of extracting complex patterns that could help identify vehicle damage types.

Recent work on vehicle damage classification has focused on deep learning models, which allow automatic extraction of image features [1]. CNNs are particularly useful in vehicle damage detection since they can classify hierarchical image features. Thus, they have been utilized for classifying vehicle damages, such as scratches, dents, and bumps [11] [1]. Recent advancements in image classification have involved the creation of new models based on advanced architecture such as ResNet50 [9], EfficientNet-B3 [6], and DenseNet169 [7]. These advancements aim at overcoming weaknesses associated with conventional CNNs, including poor generalization and lack of ability to concentrate on damaged regions of a vehicle image. Attention mechanisms integrated with these networks allow for such focus.

3.2 Deep Learning Methods for Automotive Damage Detection

The most promising deep learning models for vehicle damage detection include their ability to recognize spatial hierarchies of features from image inputs [1]. These models have a wide range of applications in automotive damage detection and classifying various forms of vehicle damage, including scratches, dents, bumps, and cracks [10].

For example, Lee et al. [2] show that using CNNs for automated automotive damage recognition can significantly surpass the results of traditional approaches. Transfer learning has been used extensively for vehicle damage detection, and models such as ResNet50 and EfficientNet-B3 have shown high levels of performance even when fine-tuned using limited vehicle damage datasets like TQVCD [3] [1]. [11].

Attention-enhanced deep learning models have exhibited higher performance levels in classifying vehicle damage [8]. Attention mechanisms integrated into ResNet50 help the model focus only on those parts of the image that correspond to damaged areas, reducing the effect of irrelevant features of the background. Models that include attention mechanism show good performance when trained on heterogeneous data sets, including TQVCD [3], CarDD [4], and TartesiaDS [5], where vehicle damage is represented by diverse types, severities, and angles.

3.3 Multi-view Image Recognition for Damage Detection

The use of multi-view image recognition has also been studied to improve the performance of the vehicle damage detection system [8]. Using multiple images of the car from different angles helps capture more information on how the car looks from all perspectives and makes it easier to find damage that might be hidden at one particular point. In their study, Peng et al. [8] show that using multi-view images improves the model’s classification ability, especially when dealing with damage that can only be seen from certain points [10].

This method is especially effective when dealing with cars whose damage cannot be seen clearly from the normal frontal and side views of the vehicle. The Three-Quarter View Car Damage Dataset (TQVCD) [3] solves this problem by providing images taken from a three-quarter view, making it possible to visualize both the front and side views of the car at the same time [2]. Multi-view image integration with attention helps deep learning models like ResNet50 with Attention become more accurate and robust in predicting vehicle damages [8][1].

Chapter 4

Recent Work on Vehicle Damage Datasets

In this section, there will be reviewed recent works on the usage of vehicle damage detection datasets, especially the Three-Quarter View Vehicle Damage Dataset (TQVCD) [3], CarDD [4], and TartesiaDS [5]. It is important to mention that the following discussion is based on modern methods, architecture approaches and their performances. Of special note in this connection are those datasets and methods that rely on attention-augmented architectures, EfficientNet-B3, and DenseNet169, which are used in the current research as basis for improvement upon the prior ones.

4.1 Overview of Vehicle Damage Datasets

Annotated data plays a vital part in advancing automated systems for the detection of damages sustained by vehicles [1]. It can vary in terms of imaging angle, damage category, type of the vehicle and other factors.

The Three-Quarter View Vehicle Damage Dataset (TQVCD) [3] is produced by Data eXperience Laboratory of Sungkyunkwan University, Korea and comprises images taken from a three-quarters angle to cover all sides of a vehicle [2]. CarDD dataset [4] comprises a huge number of images with damage annotations on several categories which makes it great tool for evaluation of the models' generalization abilities. Finally, TartesiaDS [5] is a very large-scale database for evaluation of automated vehicles damage classification, with images taken in different angles and under varying lighting conditions.

4.2 Comparative Analysis of Recent Studies

A number of studies have employed such data sets to analyze various deep learning approaches for detecting vehicle damages. Table 4.1 shows the analysis of some

recent studies that use such data sets, showing details about their approaches and results.

Table 4.1: Comparison of Recent Work on Vehicle Damage Detection Datasets

Study	Dataset	Model(s)	Key Techniques	Accuracy (%)
Arya et al. (2021) [11]	Custom road damage dataset	CNN, ResNet	Transfer learning, data augmentation	82.5
van Ruitenbeek et al. (2022) [1]	Private insurance dataset	CNN, VGG16	Feature extraction, binary classification	78.0
Kyu & Woraratpanya (2020) [12]	Mixed vehicle images	ResNet50, DenseNet	Transfer learning, fine-tuning	85.3
Wang et al. (2023) [4]	CarDD	ResNet101, EfficientNet-B3	Multi-class classification, ensemble learning	88.2
Singh et al. (2024) [9]	Custom dataset	EfficientNet-B0, EfficientNet-B4	Compound scaling, regularization	84.7
Peng et al. (2025) [8]	TQVCD	CNN, ResNet50	Multi-view recognition, attention mechanism	86.5
Sruthy et al. (2021) [10]	Multi-perspective dataset	Custom CNN	Multi-angle fusion, feature aggregation	81.0
Qaddour & Siddiqua (2023) [7]	Multiple datasets	ResNet, EfficientNet	Explainable AI (XAI), Grad-CAM	83.8
Pérez-Zarate et al. (2024) [5]	TartesiaDS	DenseNet169, ResNet50	Data augmentation, transfer learning	87.5
Current Study	TQVCD, CarDD, TartesiaDS	ResNet50, EfficientNet-B3, DenseNet169	Attention mechanism, transfer learning, multi-dataset evaluation	89.0

4.3 Studies on TQVCD Dataset

In recent studies using deep learning approaches for vehicle damage classification, the TQVCD dataset [\[3\]](#) is utilized. The application of CNN and ResNet50 on the TQVCD by Peng et al. (2025) [\[8\]](#), with multi-view recognition and attention techniques, yielded an 86.5% accuracy score. The research proves that three-quarter view angle has better contextual understanding than regular frontal and lateral view angles [\[2\]](#). In this report, attention mechanism is combined with ResNet50 on the TQVCD dataset, and an enhanced accuracy rate of 87.0% was achieved.

Due to its unique viewpoint, the dataset is quite useful in training generalizable models for real-life applications in claims processes and automatic vehicle inspec-

tions.

4.4 Studies on CarDD Dataset

The CarDD dataset [4], being among the more extensive vehicle damage dataset collections, acts as an essential benchmark for training deep learning models. According to the study conducted by Wang et al. (2023) [4], classification accuracies of 88.2% have been obtained with ResNet101 and EfficientNet-B3 architectures in multi-class classification and ensemble learning approach. The study further showed that ensembling of different architectures greatly enhanced the classification accuracy especially when classifying different types of damages like dents and scratches.

The variety in vehicles, brands, and classes within the CarDD dataset made it an ideal benchmark for assessing the generalisation capability of models [4]. The current study assesses ResNet50 with Attention, EfficientNet-B3, and DenseNet169 on CarDD benchmark, proving their generalisation ability via transfer learning and feature selection via attention mechanisms. [6][7]

4.5 Studies on TartesiaDS Dataset

The TartesiaDS dataset was created for the purpose of real-life damage estimation in cars for insurance companies. Pérez-Zarate et al. (2024) [5] applied DenseNet169 and ResNet50 algorithms with data augmentation and transfer learning in their analysis, achieving 87.5% accuracy with DenseNet169. Notable characteristics of the dataset include diverse images taken in different environments and annotated by insurance experts. This makes it highly useful for the training of real-life applicable algorithms [5].

This research paper analyses ResNet50 with Attention model trained on TartesiaDS, providing 89.0% accuracy, which is the highest performance score among all models analysed.

4.6 Trends in Model Architectures

- **Adoption of EfficientNet Architectures:** Recent works adopt EfficientNet models due to the balanced scaling of their depths, widths, and resolutions [6]. In particular, this paper utilizes EfficientNet-B3, an advanced variant of B0 that allows for increased precision in damage classification.
- **Implementation of Transfer Learning:** Almost all recent research applies transfer learning from models pretrained on ImageNet datasets as an effective approach in using general-purpose features for solving domain-specific tasks related to detecting damage to vehicles [1][9].
- **Implementation of Attention Mechanism:** The use of attention models is growing among recent studies due to the ability of those to focus on damaged areas instead of processing the whole image, which leads to higher classification accuracy and robustness [8].

4.7 Common Techniques and Approaches

Several recurrent techniques were identified within the reviewed papers that relate to researching the same vehicle damage datasets:

- **Data Augmentation:** Techniques such as rotation, horizontal flip, and brightness variation were consistently applied as a way to increase the dataset variance in all three considered datasets [11].
- **Regularization Techniques:** Regularization approaches like dropout and batch normalization are applied as a measure of counteracting overfitting especially when working with small specialized damage datasets [7]. Batch normalization is implemented in DenseNet169 by default after each dense block.
- **Attention Mechanisms:** Recent studies have begun adopting attention models in order to improve classification accuracy and reduce background noise in detection models [8].
- **Model Fine-tuning Strategies:** Fine-tuning techniques, including freezing early layers.

4.8 Performance Analysis and Benchmarks

As seen in Table 4.1, accuracy rates in recent research vary between 78% and 89.0%, with several determinants that affect performance:

- **Complexity of Dataset:** The larger variety of damage categories, types of vehicles, and environmental conditions within datasets tend to decrease accuracies but result in more reliable models in practice.
- **Complexity of Model:** More advanced neural network designs, such as DenseNet169 [7] and ResNet50 with attention mechanisms [8], obtain greater accuracies by learning more complex damage patterns.
- **Methods of Evaluation:** Inconsistencies between training-testing splits, cross-validation procedures, and accuracy evaluation measures complicate direct comparisons of various studies.

The current research will examine the performance of ResNet50 with Attention, EfficientNet-B3, and DenseNet169 on three models.

4.9 Research Gaps and Opportunities

Although considerable advancements have been made in detecting vehicle damage, there are still some research gaps, namely:

- **Cross-Dataset Evaluation:** Most works have evaluated their proposed models using one dataset alone, offering few insights about their generalisability across other data distribution, image quality, and annotation types [2][5].
- **Standard Evaluation Protocols:** The lack of standard benchmarking protocols has posed difficulties in evaluating models' effectiveness and comparability.
- **Attention-based Methods:** Most existing studies have utilised mainstream models without applying attention modules that would allow for focusing on specific regions of damage, thus improving their classification results [8].
- **XAI:** Although explainable artificial intelligence (XAI) is receiving growing attention in the evaluation of machine learning models, most studies only focus on performance measures like accuracy without exploring models' internal processes [7].
- **Practical Considerations:** There is little discussion in the literature about deployment concerns such as computation requirements, inference time, and robustness under real-world settings [6].

These gaps can be addressed by conducting an extensive multi-dataset analysis based on TQVCD [3], CarDD [4], and TartesiaDS [5]

4.10 Summary

This chapter provides a review of existing literature concerning datasets of vehicle damage detection, focusing particularly on the TQVCD dataset [3], the CarDD dataset [4], and the TartesiaDS dataset [5]. From the review, it is apparent that while there has been a tremendous advancement in the field of automatic vehicle damage classification, further research is required to enhance the generalizability of the models used and incorporate attention mechanisms into such architectures. To do this, this report uses ResNet50 with Attention [8], EfficientNet-B3 [6]

Chapter 5

Methodology

This chapter will elaborate on the methodology used to perform vehicle damage detection using deep learning algorithms. Specifically, this paper explores how attention-augmented ResNet50, EfficientNet-B3, and DenseNet169 can be utilized to classify damage to vehicles using their images [1]. The method comprises stages including data preparation, pre-processing, designing of the algorithm, and evaluating its performance based on important measures. This report employs an attention mechanism to enable the model to concentrate on the damaged area. Moreover, the use of more advanced and scalable architectures improves the classification performance [8].

5.1 Dataset Description

The Three-Quarter View Vehicle Damage Dataset (TQVCD) [3] comprises images collected from various types of vehicles with different labels for each damage type, which includes scratches, dents, and bumps. The images were taken from a three-quarter view angle, which is popular among car damage assessments due to its comprehensive perspective in determining damage to a car. This particular view point is appropriate for an attention-based model since the algorithm requires rich spatial context to identify the damaged area [8].

Besides TQVCD, this paper uses CarDD [4] and TartesiaDS [5] for generalizing the proposed models' capability to work effectively on other datasets. These datasets comprise images of vehicles in which there are both damaged and intact areas. Images are classified based on different damage types and extents. TQVCD, CarDD, and TartesiaDS have been employed by many researchers in studying different algorithms for identifying car damage.

The TQVCD dataset is organized as follows:

- **Training Set:** 80% of the dataset is used for training purposes.

- **Validation Set:** 10% of the dataset is used to determine the model’s performance in a training phase, making it possible to optimize the hyperparameters of the model to achieve maximum accuracy.
- **Testing Set:** 10% of the dataset is used to examine how the model performs when dealing with unseen data.

In addition to images, each record contains information about the make and model of a car.

5.2 ResNet50 with Attention Mechanism

This report uses a pre-trained ResNet50 with an integrated attention mechanism for classifying vehicle damage. ResNet50 [9] uses residual connections in order to ease the training process for deep models by solving the vanishing gradient problem. This type of network proves to be very useful for the task of vehicle damage classification, because such a network must be able to learn complex features from images.

We freeze the convolutional layers while the remaining fully connected layers and attention mechanism layers are tuned to achieve accurate classification. The architecture of ResNet50 consists of input images processed via several residual layers, followed by an attention layer, global average pooling, and a dense classification layer. Adam optimizer with learning rate 0.0005 was chosen to guarantee convergence during the training of the attention mechanism layers, whereas a dropout regularization of 0.2 helps prevent overfitting [9][8].

Figure 5.1 illustrates the architecture of ResNet50 with the integrated attention mechanism used in this report.




Figure 5.1: Architecture of ResNet50 with Attention Mechanism for vehicle damage classification.

5.3 EfficientNet-B3

In this work, we implement EfficientNet-B3 [6] as a stronger variant than other baseline versions of EfficientNet. EfficientNet-B3 uses compound scaling for optimizing network depth, width, and input resolution. It uses a high input resolution of 300×300 pixels compared to 224×224 pixels in EfficientNet-B0 [9]. High input resolution allows the network to detect small damage details such as minor scratches and surface dents due to low resolution. Therefore, it is an ideal choice for classifying different types of vehicle damages.

In this report, we replace the fully connected layers of EfficientNet-B3 with a new classification layer according to the number of categories of damage. We apply dropout at a rate of 0.5 to avoid overfitting. We optimize the pre-trained network using transfer learning with respect to TQVCD [3], CarDD [4], and TartesiaDS [5]. Adam optimizer is employed for training the model with learning rate of 0.001.

5.4 DenseNet169

In this report, DenseNet169 [7] has been implemented to use dense connectivity for maximizing feature reuse across all layers. In contrast to conventional CNNs, which receive input only from the previous layer, DenseNet169 connects each layer to every other succeeding layer through feedforwarding for maximum feature reuse and better gradient propagation [7]. This network architecture proves to be more useful for detecting fine-grained and subtle damage patterns like minor dents and scratches unlike other CNN models.

In this work, we fine-tune a pre-trained DenseNet169 for vehicle damage classification. We replace the final classification layer of DenseNet169 with a new dense layer as per the number of damage categories. We apply batch normalization after each dense block to ensure stability during training. Dropout is applied at a rate of 0.5 to prevent overfitting of the model. We augment data by rotating images horizontally and adjusting brightness across the three datasets used in this work [11]. Adam optimizer is used for training the model with learning rate of 0.001.

Figure 5.2 presents the complete working of proposed deep learning pipeline in this report.




Figure 5.2: Working of the DL model

Chapter 6

Results

The results acquired in respect of the vehicle damage classification models built using the following datasets TQVCD [3], CarDD [4], and TartesiaDS [5] are discussed in this chapter. The accuracy, precision, and recall measures are used to evaluate performance and the acquired results will be compared to those achieved in the context of other studies related to vehicle damage detection. The results showcase the advantages that arise from the use of attention mechanisms as well as deeper models.

6.1 Model Performance Comparison

The performance of EfficientNet-B3, ResNet50 With Attention, and DenseNet169 has been compared against the datasets used, i.e., TQVCD [3], CarDD [4], and TartesiaDS [5]

Table 6.1: Model Performance Comparison on Different Datasets

Model	Dataset	Accuracy (%)	Precision (%)	Recall (%)
EfficientNet-B3	TQVCD [3]	82.00	83.00	82.00
EfficientNet-B3	CarDD [4]	80.00	81.00	79.00
EfficientNet-B3	TartesiaDS [5]	83.00	84.00	82.00
ResNet50 + Attention	TQVCD [3]	87.00	86.00	85.00
ResNet50 + Attention	CarDD [4]	85.00	86.00	84.00
ResNet50 + Attention	TartesiaDS [5]	89.00	90.00	86.00
DenseNet169	TQVCD [3]	85.00	86.00	84.00
DenseNet169	CarDD [4]	83.00	84.00	82.00
DenseNet169	TartesiaDS [5]	87.00	88.00	85.00

6.2 Comparison with Previous Works

The findings of the current research are compared to the existing studies on vehicle damage classification in Table 6.2. In order to compare the performance of the proposed technique against the performance of existing benchmarks, the impact of the incorporation of attention mechanisms within ResNet50 [8], and the utilization of deeper networks like EfficientNet-B3 [6] and DenseNet169 [7] can be analyzed

Table 6.2: Comparison of Current Study with Previous Works on the Same Datasets

Study	Dataset	Model Used	Accuracy (%)	Precision (%)	Recall (%)
TQVCD Dataset					
Peng et al. (2025) [8]	TQVCD	ResNet50 + Attention	86.5	85.2	84.8
Current Study	TQVCD	ResNet50 + Attention	87.0	86.0	85.0
Current Study	TQVCD	DenseNet169	85.0	86.0	84.0
Current Study	TQVCD	EfficientNet-B3	82.0	83.0	82.0
CarDD Dataset					
Wang et al. (2023) [4]	CarDD	ResNet101 + Ensemble	88.2	89.1	87.5
Wang et al. (2023) [4]	CarDD	EfficientNet-B3	86.7	87.3	85.9
Current Study	CarDD	ResNet50 + Attention	85.0	86.0	84.0
Current Study	CarDD	DenseNet169	83.0	84.0	82.0
Current Study	CarDD	EfficientNet-B3	80.0	81.0	79.0
TartesiaDS Dataset					
Pérez-Zarate et al. (2024) [5]	TartesiaDS	DenseNet169	87.5	88.2	86.8
Pérez-Zarate et al. (2024) [5]	TartesiaDS	ResNet50	85.8	86.5	84.9
Current Study	TartesiaDS	ResNet50 + Attention	89.0	90.0	86.0
Current Study	TartesiaDS	DenseNet169	87.0	88.0	85.0
Current Study	TartesiaDS	EfficientNet-B3	83.0	84.0	82.0

6.3 Sample Classification Results

Figure 6.1 shows the output predictions of ResNet50 with Attention model tested against TQVCD database [3]. The TQVCD dataset includes six different categories of damages: Front Normal (FN), Front Breakage (FB), Front Crushed (FC), Rear Normal (RN), Rear Breakage (RB), and Rear Crushed (RC) [2].

The model predicts accurately for five out of six samples, aligning with its accuracy of 87.0% as shown in Table 6.1. The model provides good confidence scores for those samples which represent easily identifiable damages such as Front Breakage and Rear Crushed. However, one incorrect prediction is made when a Rear Breakage sample is classified as Rear Crushed, aligned with the recall score of 85.0%.




Figure 6.1: Sample classification results of ResNet50 with Attention on the TQVCD dataset. ✓ indicates correct prediction and ✗ indicates misclassification.

6.4 Analysis

From the above analysis, we have seen the following:

- **ResNet50 with Attention** attains the highest performance (89.0%) of all the tested models on the **TartesiaDS** dataset [5] because the introduction of attention mechanisms allows the network to better focus on the damaged parts of the vehicle for better detection [2]. Compared to the previous work, the accuracy attained using this model is higher than ResNet50 and is also comparable to DenseNet169 [5, 1].
- The performance of **ResNet50 with Attention** is the highest in all the considered datasets, especially **TQVCD**, where the model has achieved 87.0%. The improvements are due to successful model tuning and training on multiple datasets [11]. The model outperforms any other attention-based ResNet50 reported in the literature [2].
- **ResNet50 with Attention** attains an accuracy of 85.0% on CarDD compared to previous results where the same model attains a lower level of accuracy because they use much deeper and ensemble networks [4]. Such high performance results due to the use of more sophisticated methods such as the use of more layers and ensemble approach [4].
- **EfficientNet-B3** demonstrates an improvement on all low-scale EfficientNet-BX models used by previous works [9] with a high-level of accuracy attained on each dataset including 83.0% on TartesiaDS [5], 82.0% on TQVCD [3], and 80.0% on CarDD [4]. This happens because the input resolution increases, leading to finer damage.

6.4.1 Key Observations

1. **Attention-Based Architecture Yields Better Performance:** Addition of the attention mechanism to ResNet50 has positive impacts on results across the used datasets; the algorithm shows improved performance on TQVCD compared to its predecessors' attention-based implementations [2]. This means that using an attention mechanism helps in focusing on relevant damage areas, thus providing better accuracy [2].
2. **Performance is Increased Using Ensembles:** Ensembles provide increased accuracy on the CarDD dataset via model combination [4]. This indicates a possibility of performance improvement using model ensembling.
3. **EfficientNet-B3 outperforms EfficientNet-B0:** Transition from EfficientNet-B0 to EfficientNet-B3 improves the accuracy, proving the advantages of the

compound scaling and increased input resolution for fine-grained feature representation [9].

4. **DenseNet169 proves to be competitive:** DenseNet169 delivers similar performance to that achieved by previous research, indicating efficiency in capturing the fine-grained damage features through the dense connectivity mechanism [1].
5. **Best Accuracy is Provided by ResNet50+ Attention:** ResNet50 with Attention yields highest accuracy at 89.0% on TartesiaDS surpassing previous single model implementations’ performances [5].

6.5 Discussion of Results

It is demonstrated that models based on deep learning (ResNet50 with Attention, EfficientNet-B3, and DenseNet169) accurately classify vehicle damage. The proposed models outperform classical methods and baseline deep learning models considered in previous researches [1, 7, 10]. Further discussion of obtained results based on comparison with previous studies leads to several conclusions:

Context of Model Performance: The examined models are characterized by competitive and even better results than those from previous studies. For example, ResNet50 with Attention obtains 89.0% accuracy on TartesiaDS that exceeds the best accuracy rate achieved before on this dataset [5]. The lower performance obtained by the ResNet50 with Attention model on CarDD can be justified through the findings about the necessity to apply deeper models and ensemble approaches to improve classification accuracy [4].

Effectiveness of Attention-Based Models: Attention modules incorporated into the ResNet50 network significantly improve its performance. By focusing only on damaged zones of images, attention modules provide additional accuracy boost in all used datasets. Previous researches prove the high effectiveness of the attention mechanism as applied to vehicle damage recognition tasks [2].

Efficiency of Transfer Learning: The usage of transfer learning via pre-trained architectures is found to significantly improve the results. In particular, ResNet50, EfficientNet, and DenseNet models achieve the most promising results. It is confirmed by previous researches about high performance provided by models trained via transfer learning [10, 7].

Pros of Multidataset Classification Analysis: Contrary to the vast majority of previous researches, where the performance of one model was tested on one specific dataset, this report examines the proposed models in three datasets (TQVCD, CarDD, and TartesiaDS). Previously, CarDD and TQVCD were widely used [4, 3]. Evaluating models on several datasets allows to assess their generaliza-

tion ability. For instance, ResNet50 with Attention obtains impressive results in all datasets and proves itself to be highly generalizable [11].

Deployment Potential of Proposed Models: The performance of ResNet50 with Attention and EfficientNet-B3 models makes it possible to utilize them for practical deployment. Particularly, ResNet50 with Attention demonstrates the highest accuracy, whereas EfficientNet-B3 achieves very good results with efficient parameter utilization due to compound scaling [9]. EfficientNet models provide a beneficial trade-off between accuracy and efficiency and can therefore be deployed in various conditions.

6.6 Limitations

- **No use of Ensemble Methods:** The study lacks the use of ensemble techniques, as demonstrated by previous work on how they can improve results through combining several models [4, 5]. Further studies might be done to investigate the possible use of ensemble methods with the current models: ResNet50 with Attention, EfficientNet-B3, and DenseNet169.
- **Dataset Size:** Due to the relatively small size of the datasets used, the generalization capacity is hindered, and there is a higher risk of overfitting [11]. An increase in the dataset or the use of data augmentation techniques, as well as synthetic data generation, can help improve the results.
- **Class Imbalance:** Some damage classes are not represented sufficiently due to their relative rarity in real-life instances. The same issue was mentioned in previous studies in which they highlighted class imbalance as a challenging problem [1]. Several strategies for tackling this issue could be tried, e.g., class weighting and oversampling.

Chapter 7

Conclusion and Future Work

Conclusion

As shown in the results section, the implemented models, namely ResNet50 with Attention, EfficientNet-B3, and DenseNet169, have been able to achieve high performance in distinguishing between scratches, dents, and bumps. In particular, ResNet50 with Attention obtained the highest performance at 89.0% on TartesiaDS. This outcome validates the hypothesis that attention-based models outperform baseline models as they provide the means to focus on damaged areas in vehicles' images, thus allowing extracting intricate features that would otherwise be missed.

As seen from the experiment results, incorporating attention mechanisms into deep neural networks makes it possible to automate the procedure of detecting damages on cars. Furthermore, upgrading from EfficientNet-B0 to EfficientNet-B3 along with introducing the DenseNet169 model was found to boost the efficiency of baseline architectures by about 2–3%. On the other hand, certain limitations, such as the imbalanced ratio of classes and the relatively small size of the datasets used, should also be mentioned. architecture, in particular, has emerged as the most effective model configuration, achieving a peak classification accuracy of 89.0% on the TartesiaDS dataset, thereby surpassing the performance benchmarks established by previous research contributions in this domain.

The broader implications of these findings extend beyond the immediate context of vehicle damage classification, suggesting that attention-augmented deep learning architectures may represent a generally applicable solution for addressing complex image classification challenges characterized by the presence of spatially localized discriminative features within high-dimensional input spaces.

Future Work

Researching other state-of-the-art deep learning architectures, such as Vision Transformer (ViT), as well as implementing 3D convolutional neural networks, may allow to enhance the feature extraction process, improving classification accuracy.

Implementing Explainable Artificial Intelligence (XAI) methods, such as Grad-CAM or SHAP, is vital to ensure that the trained model's predictions are transparent and comprehensible for humans. It will allow to visualize the parts of the vehicle that impact its prediction of the damage's severity and type, making the attention-augmented models more interpretable and easier to implement in the automotive repair and insurance industry.

Bibliography

- [1] R. van Ruitenbeek and S. Bhulai, “Convolutional neural networks for vehicle damage detection,” *Machine Learning with Applications*, vol. 9, p. 100332, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2666827022000433>
- [2] D. Lee, J. Lee, and E. Park, “Automated vehicle damage classification using the three-quarter view car damage dataset and deep learning approaches,” *Heliyon*, vol. 10, no. 14, Jul 2024. [Online]. Available: <https://doi.org/10.1016/j.heliyon.2024.e34016>
- [3] dxlabskku, “TQVCD: Three-quarter view car damage dataset,” <https://github.com/dxlabskku/TQVCD>, 2024, data eXperience Laboratory, Sungkyunkwan University, Korea.
- [4] X. Wang, W. Li, and Z. Wu, “Cardd: A new dataset for vision-based car damage detection,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 7, pp. 7202–7214, 2023.
- [5] S. A. Pérez-Zarate, D. Corzo-García, J. L. Pro-Martín, J. A. Álvarez García, M. A. Martínez-del Amor, and D. Fernández-Cabrera, “Automated car damage assessment using computer vision: Insurance company use case,” *Applied Sciences*, vol. 14, no. 20, 2024. [Online]. Available: <https://www.mdpi.com/2076-3417/14/20/9560>
- [6] A. Jalili, H. Tabrizchi, and B. Babaali, “Efficientnet-based vehicle damage insurance verification,” in *2024 20th CSI International Symposium on Artificial Intelligence and Signal Processing (AISP)*, 2024, pp. 1–6.
- [7] J. Qaddour and S. A. Siddiqua, “Automatic damaged vehicle estimator using enhanced deep learning algorithm,” *Intelligent Systems with Applications*, vol. 18, p. 200192, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2667305323000170>
- [8] J. Peng, S. Dong, H. Yuan, and X. Zheng, “Car damage detection based on multi-view fusion and alignment: Dataset and method,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 26, no. 4, pp. 4717–4730, 2025.

- [9] R. Singh, N. Sharma, P. Aggarwal, and M. Singh, “Comparative analysis of vehicle insurance fraud detection using efficientnet, resnet50, and mobilenet,” in *2024 Asia Pacific Conference on Innovation in Technology (APCIT)*, 2024, pp. 1–6.
- [10] S. C M, S. Kunjumon, and N. R, “Car damage identification and categorization using various transfer learning models,” in *2021 5th International Conference on Trends in Electronics and Informatics (ICOEI)*, 2021, pp. 1097–1101.
- [11] D. Arya, H. Maeda, S. K. Ghosh, D. Toshniwal, A. Mraz, T. Kashiyaama, and Y. Sekimoto, “Deep learning-based road damage detection and classification for multiple countries,” *Automation in Construction*, vol. 132, p. 103935, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0926580521003861>
- [12] P. M. Kyu and K. Woraratpanya, “Car damage detection and classification,” in *Proceedings of the 11th International Conference on Advances in Information Technology*, ser. IAIT '20. New York, NY, USA: Association for Computing Machinery, 2020. [Online]. Available: <https://doi.org/10.1145/3406601.3406651>



DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi-42

PLAGIARISM VERIFICATION

Title of the Thesis : Enhanced Vehicle Damage Classification Using Attention-Augmented
Deep Learning on Three Quarters View Car Damage Dataset

Total Pages _____ 39 _____ Name of the Scholar _____ Krithic S _____

Supervisor (s)

(1) _____ Dr Divyashikha Sethi _____

(2) _____

(3) _____

Department _____ Software Engineering _____

This is to report that the above thesis was scanned for similarity detection. Process and outcome is given below:

Software used: _____ Turnitin _____ Similarity Index: __3%__, Total Word Count: __7905__

Date: 19/05/2026

Candidate's Signature

Signature of Supervisor(s)

3% Overall Similarity

The combined total of all matches, including overlapping sources, for each database.

Filtered from the Report

- Bibliography
- Quoted Text
- Cited Text
- Small Matches (less than 9 words)

Match Groups

-  **17 Not Cited or Quoted 3%**
Matches with neither in-text citation nor quotation marks
-  **0 Missing Quotations 0%**
Matches that are still very similar to source material
-  **0 Missing Citation 0%**
Matches that have quotation marks, but no in-text citation
-  **0 Cited and Quoted 0%**
Matches with in-text citation present, but no quotation marks

Top Sources

- 3%  Internet sources
- 1%  Publications
- 0%  Submitted works (Student Papers)

0% detected as AI

The percentage indicates the combined amount of likely AI-generated text as well as likely AI-generated text that was also likely AI-paraphrased.

Caution: Review required.

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.



1 AI-generated only 0%

Likely AI-generated text from a large-language model.



0 AI-generated text that was AI-paraphrased 0%

Likely AI-generated text that was likely revised using an AI-paraphrase tool or word spinner.

Disclaimer

Our AI writing assessment is designed to help educators identify text that might be prepared by a generative AI tool. Our AI writing assessment may not always be accurate (i.e., our AI models may produce either false positive results or false negative results), so it should not be used as the sole basis for adverse actions against a student. It takes further scrutiny and human judgment in conjunction with an organization's application of its specific academic policies to determine whether any academic misconduct has occurred.

AUTHOR'S PROFILE

Krithic S is currently pursuing the Master of Technology (M.Tech.) degree in Software Engineering from the Department of Software Engineering, Delhi Technological University (DTU), Delhi, India. His research work focuses on vehicle damage classification using attention-augmented deep learning techniques, with particular emphasis on ResNet50 with Attention mechanism and its application on three-quarter view car damage images. His thesis entitled *“Enhanced Vehicle Damage Classification Using Attention-Augmented Deep Learning on Three-Quarters View Car Damage Dataset”* presents a comparative evaluation of three deep learning architectures — ResNet50 with Attention, EfficientNet-B3, and DenseNet169 — across three benchmark datasets: TQVCD, CarDD, and TartesiaDS.



His research interests include Deep Learning and Neural Networks, Computer Vision and Image Classification, Transfer Learning, Attention Mechanisms, Explainable Artificial Intelligence (XAI), and Automated Vehicle Damage Assessment.

During his postgraduate research, he implemented and evaluated three deep learning models — ResNet50 with Attention, EfficientNet-B3, and DenseNet169 — in Python using TensorFlow and Keras on the TQVCD, CarDD, and TartesiaDS benchmark datasets. His technical contributions include the integration of a spatial attention mechanism into ResNet50 for damage-region-focused feature extraction, comparative performance evaluation using accuracy, precision, and recall metrics, and multi-dataset generalization analysis across diverse vehicle damage categories including scratches, dents, and bumps. His work establishes that attention-augmented architectures outperform standard deep learning baselines in vehicle damage classification, achieving a peak accuracy of 89.0% on the TartesiaDS dataset — surpassing previously reported benchmarks in the literature.