International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

AUTOMATIC DAMAGED VEHICLE ESTIMATOR USING DEEP LEARNING ALGORITHMS

Priyansh Verma[1], Kaustubh Shankar[2], Rahul Biswas[3], Prof. B. K. Singh[4]

2021UCS1700
2021UCS1669
2021UCS1666
[4] Under the Supervision Of
Department Of Computer Science & Engineering Netaji Subhas University of Technology (NSUT) New Delhi, India-110078

ABSTRACT :

Manual car damage evaluations, which are sometimes laborious, uneven, and prone to human mistake, are ongoing issues for the auto insurance sector. This project uses deep learning approaches to automate the damage evaluation process in an AI-driven manner. Utilizing a pre-trained ResNet50 model that has been improved by transfer learning and adding Generative Adversarial Networks (GANs) to augment data, the system efficiently identifies damage, categorizes its level of severity, and standardizes the evaluation process. ResNet50's usage of skip connections allows for more reliable feature extraction, and GANs help overcome data availability constraints, improving generalization. Insurers can make more consistent and trustworthy decisions thanks to this automated approach, which also speeds claim processing and lessens subjectivity.Moreover, the ability of the system to consistently and objectively assess damage ensures that insurance companies make more reliable financial decisions. Through improvements in accuracy and efficiency, the study shows how deep learning may revolutionize insurance operations.

Chapter 1 INTRODUCTION

Background

In today's rapidly evolving technological landscape, industries worldwide are embracing automation to boost productivity and reduce errors. The automobile insurance sector stands to benefit significantly from these advancements, as precise vehicle damage detection and classification are crucial for efficient claim processing. Historically, manual assessment processes have been labor-intensive, prone to fraud, and subject to errors, leading to substantial financial losses due to claim leakage. This not only affects insurance companies but also impacts policyholders who rely on timely and fair settlements.

Deep Learning (DL) and Artificial Intelligence (AI) present encouraging answers to these problems. The development of automated systems that assess car damage, measure its extent, and produce accurate repair cost estimates has been made possible by deep learning's success in tasks like object identification, image segmentation, and image recognition. In the end, both insurers and clients stand to gain from this automation's ability to cut financial losses, improve decision-making, and lessen the need for people.

Specifically, ResNet (Residual Network) architectures can leverage their ability to train deeper networks for improved accuracy in detecting complex features, while Generative Adversarial Networks (GANs) can generate realistic images of damaged vehicles, enhancing training datasets and improving the estimator's robustness. By integrating ResNet and GANs, the industry can automate vehicle damage assessment, streamlining claim processing, enhancing accuracy, and accelerating approvals. This integration boosts customer satisfaction and supports predictive analytics for better risk management, aligning operations with evolving insurer and customer needs.

Moreover, this automation brings about a more personalized experience for policyholders. With faster and more accurate claims processing, customers can expect quicker resolutions and more transparent communication throughout the process. This not only enhances customer satisfaction but also fosters trust between insurers and their clients. As technology continues to evolve, embracing AI and deep learning will be crucial for insurance companies to stay competitive and deliver the high-quality service that customers expect in today's digital age.

Motivation

Insurance companies and policyholders alike frequently find the handling of auto insurance claims to be a tedious, error-prone, and frustrating procedure. When a car is damaged, the duration taken by an inspector to evaluate the damage by hand, calculate the cost of repairs, and process a claim can seem drawn out and ineffective. In addition to being sluggish, this manual method is prone to errors and discrepancies, which may result in insurers paying out too much or too little. Insurance firms lose millions of rupees annually due to this problem, which is referred to as claim leakage, straining their bottom line and client relations.

Automating the claims process is facilitated by the development of AI-driven models for solving such issues. By utilising these models, insurance companies can expedite the process of assessing damage, enabling real-time evaluations that result in a significant reduction in the time lag between the submission of claims and their resolution. Artificial intelligence (AI) systems have the ability to evaluate damage photographs, categorise the level of repairs required, and generate precise cost estimates while reducing the possibility of human error.

The transition towards automation reduces claim leakage by streamlining operations and improving the accuracy of damage assessments. Enhancing the consistency and dependability of assessments enables insurers to more efficiently manage expenses and distribute resources, which eventually results in superior financial performance. Additionally, by providing quicker, more transparent processing of claims, this technology-driven strategy strengthens bonds with policyholders and raises consumer satisfaction and loyalty. Adopting these technologies will be essential for insurance businesses hoping to stay competitive in a continuously changing market as the sector continues to evolve.

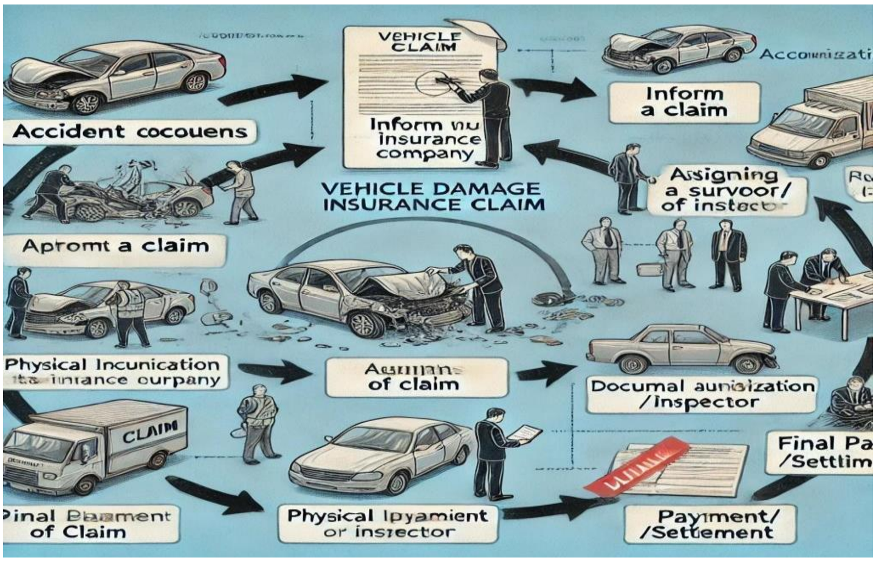

Figure 1. The manual process of insurance claim

GANs

Two neural networks—a discriminator and a generator—compete with one another to form Generative Adversarial Networks (GANs), a kind of deep learning model. While the discriminator attempts to discern between authentic and fraudulent photos, the generator produces artificial images. The generator gains the ability to create incredibly lifelike images through this adversarial process. In order to supplement the training dataset, this study uses GANs to produce artificial images of wrecked cars. By expanding the range of damage kinds, vehicle models, and environmental variables the model is exposed to during training, this helps address the issue of sparse and unbalanced real-world data.

ResNet50

The concept of residual learning is introduced by ResNet50, a 50-layer deep convolutional neural network. ResNet50 employs shortcut connections, which bypass one or more layers, in contrast to conventional CNNs. This facilitates the flow of gradients during backpropagation. This improves training accuracy and resolves the issue of disappearing gradients in very deep networks. This project uses ResNet50 to classify whether an automobile is damaged and to determine whether a car is there in a picture. It is especially useful for deciphering tiny visual clues associated with damage because of its capacity to extract and identify intricate characteristics from photos.

Transfer Learning

A deep learning technique called transfer learning involves adapting a model that was trained on a sizable, generic dataset for a new but related goal. We utilize a pre-trained model, such as ResNet50, which has previously learned to detect common image features from a dataset like ImageNet, rather than developing a model from scratch. Then, using our unique dataset of automobile photos, we refine this model. This enhances performance while drastically cutting down on training time and data needs. Transfer learning enables the model to concentrate on learning task-specific patterns, like damage severity and location, more efficiently by utilizing previously learnt characteristics.

Chapter 2

PROBLEM STATEMENT AND LITERATURE SURVEY

Problem Statement

In the insurance industry, manual systems for evaluating car damage are plagued by inaccuracies and inefficiencies. When a car is damaged, an insurance adjuster typically inspects it, assesses the damage, and determines the repair cost. However, this manual method is susceptible to subjectivity, variations in appraisal, and human error. These issues can lead to overestimation or underestimation of damage, resulting in incorrect claim payments. Additionally, the manual process is slow and time-consuming, frustrating clients and increasing operational strain on insurance companies by causing delays in claim settlements.

Besides these inefficiencies, insurance firms often struggle to accurately estimate repair costs due to the complexity of vehicle damage evaluations. Furthermore, the traditional approach fails to reliably classify the extent of damage, an essential function for estimating repair costs and making timely payments. The use of human inspectors results in assessment variability, as different adjusters may provide different estimates for the same damage, adding to the process's complexity.

This project proposes an automated method for evaluating car damage using deep learning models and AI-based tools. The system can automatically identify, segment, and evaluate car damages from images by employing ResNet (Residual Network) architectures for improved image recognition. ResNet's ability to train deeper networks enhances accuracy in detecting complex features within images. Additionally, Generative Adversarial Networks (GANs) can generate realistic images of damaged vehicles, aiding in the creation of diverse training datasets and improving the robustness of the automated damage estimator.

By automating the evaluation process, damage assessments become more accurate and consistent, while claim processing times are reduced, speeding up settlements and increasing customer satisfaction. The automated technique reduces the possibility of inconsistent evaluations by ensuring that damage assessments are more objective and standardized. Additionally, insurance companies can optimize their operational workflows by making better decisions about repair costs and offering a transparent, data-driven examination of the damage.

In conclusion, this study develops a deep learning model that automates the evaluation process to solve the inefficiencies and errors present in conventional car damage evaluations. The suggested method increases precision, shortens the time needed to process claims, and guarantees more dependable results for policyholders and insurance firms alike.

Key Challenges

1. ***Data Scarcity and Imbalance***

Obtaining a large and diverse dataset of labeled car damage images is difficult. Real-world datasets often lack variety in damage types, angles, lighting, and car models. Additionally, datasets are usually imbalanced, with more examples of undamaged vehicles than various levels of damage, which can negatively impact model training and prediction accuracy.

2. ***Complexity of Damage detection and localization***

Car damage varies greatly in location, size, and severity. Identifying subtle scratches or dents—especially under poor lighting or complex backgrounds—is a challenging task for image recognition models. Accurately segmenting the damaged regions to determine the severity and location (front, rear, side) adds another layer of difficulty.

3. ***Generalization Across Vehicle Types and Conditions***

Deep learning models may perform well on known vehicle types but struggle with new or uncommon car models, aftermarket modifications, or non-standard damage patterns. Ensuring the system generalizes well across diverse car designs, colors, and environmental conditions requires a robust and representative training set.

4. ***Integration with Real-world Insurance workflows***

Deploying the model in practical insurance claim processes involves challenges beyond model accuracy. It must interface with insurance systems, provide interpretable results, and comply with legal and privacy requirements. Resistance to adopting AI-based tools and trust issues among stakeholders can also hinder integration.

Literature Survey

With the advent of deep learning, researchers began applying Convolutional Neural Networks (CNNs) to picture recognition challenges. According to ***Patil et al. (2022),*** **[6]** the suggested framework used convolutional autoencoders to reconstruct the image for feature extraction. Because tagged datasets were difficult to find, an augmented unlabeled dataset was also employed. Using SoftMax classifiers and SVMs, the study sought to classify car damage into several categories, such as dents, collisions, mirror scratches, and so on. Using pre-trained CNN models for transfer learning, like AlexNet and VGG, demonstrated significant improvements in avoiding data overfitting.

Kyu et al. (2020). **[3]** suggested a novel approach to categorising auto damage based on location and severity of damage. suggested a method for categorising damages based on the pre-trained **CNN VGG models** itself, taking into account transfer learning during model training to prevent overfitting. They also
incorporated **L2 (Ridge) Regularization**, which might be helpful for prediction.
Using pre-trained models for damage detection and **YOLO** for damage localization, the study by ***Gandhi et al. (2021)*** **[1]** on deep learning-based on damage identification, classification, and severity evaluation offers a thorough framework for automating the insurance claims process. **DenseNet's** effectiveness in transfer learning and tuning turned out to be excellent.

Liu et al., 2022 **[4]** propose an automated vehicle damage inspection system using deep convolutional neural networks (CNNs) and multi-view image fusion to enhance detection accuracy and robustness. Their framework extracts visual features from multiple camera angles and integrates them using a view aggregation mechanism, enabling precise localization and classification of surface damage. The system outperforms single-view baselines in terms of detection accuracy and generalizability, offering a scalable solution for applications in automotive insurance, fleet management, and rental services.

Sarang Kadakia et al. (2023) **[2]** present a comprehensive analysis on leveraging deep learning for automotive damage detection. Their study employs various architectures, including VGG16, ResNet50, CNN (Sequential), Mask R CNN, SSD, and RetinaNet, to identify vehicle damage types. The results show that VGG16 achieves the highest accuracy of 97.01%, followed by ResNet50 at 94.33%, CNN (Sequential) at 92.28%, Mask RCNN at 90.00%, and SSD and RetinaNet at 85.00%. This work highlights the effectiveness of deep learning in automating damage assessment and demonstrates the potential of integrating multiple models to optimize vehicle appraisal processes, with **VGG16** exhibiting the best performance.

Mohamed et al. (2023) **[5]** explore the utilization of **Generative Adversarial Networks (GANs)** for data augmentation in car damage detection datasets. By generating synthetic images, GANs effectively mitigate data imbalance issues, thereby augmenting dataset size and enhancing class balance. This approach outperforms traditional augmentation techniques, resulting in improved accuracy and robustness of deep learning models for detecting and classifying car damages. The application of GANs demonstrates a promising strategy for overcoming data limitations and enhancing model performance in automotive damage detection tasks.

Wang et al. (2023) **[7]** explore how Vision Transformers (ViTs) can improve the accuracy of automotive damage detection, especially in cases where traditional CNNs fall short. By fine-tuning a pre-trained ViT model on vehicle damage datasets, the system is able to better recognize complex and subtle patterns like dents, scratches, and overlapping damage. The global attention mechanism of ViTs helps capture context more effectively, leading to better performance in both identifying and localizing damage. Their results show that transformers
offer a promising step forward in making automated vehicle inspection more reliable and precise.
Recent advancements in deep learning and computer vision have improved vehicle damage detection. ***Warungse et al. (2023)*** **[8]** propose a framework using Convolutional Neural Networks like ResNet to assess and locate vehicle damages, enhancing safety and efficiency. Object detection techniques such as YOLO and Faster **R-CNN** complement this by providing rapid and precise localization of damaged areas. Integrating these architectures could create a robust system, combining YOLO's speed with Faster R-CNN's precision for enhanced damage assessment.

Ref	Authors (Year)	Methods/Models	Key Contribution

[1]	Gandhi et al. (2021)	CNNs, YOLO, DenseNet	Full pipeline for damage detection, localization, severity classification.
[2]	Kadakia et al.(2023)	VGG16, ResNet50, Mask R-CNN, SSD	VGG16 highest accuracy (97.01%); multi-model comparison.
[3]	Kyu et al. (2020)	VGG + TL, L2 Regularization	Damage categorized by location/severity; regularization avoids overfitting.
[4]	Liu et al. (2022)	Multi v CNN, Fusion	Better detection via multi-angle image integration.
[5]	Mohamed et al. (2023)	GANs	Synthetic images for class balance; improved accuracy.
[6]	Patil et al. (2022)	Autoencoders, SVM, AlexNet, VGG	Used unlabeled data; categorized damage types.
[7]	Wang et al. (2023)	Vision Transformers (ViTs)	Improved subtle damage detection using global attention.
[8]	Warungse et al. (2023)	ResNet, YOLO, Faster R-CNN	Combined YOLO's speed + R-CNN's precision for robust assessment.

Table 1. Comparison of Literature Survey

Chapter 3 Methodology

Objective

This project aims to develop an AI-based system that enhances the efficiency, accuracy, and transparency of vehicle damage assessment in the automotive insurance industry. By replacing manual evaluations with deep learning-powered automation, the system seeks to transform the way claims are processed and settled.

1. ***Automated Damage Detection and Severity Classification***

Leverage a pre-trained **ResNet50** model, fine-tuned through transfer learning, to automatically detect the presence of vehicle damage, determine its location (front, rear, side), and classify its severity (minor, moderate, severe) using photographic inputs.

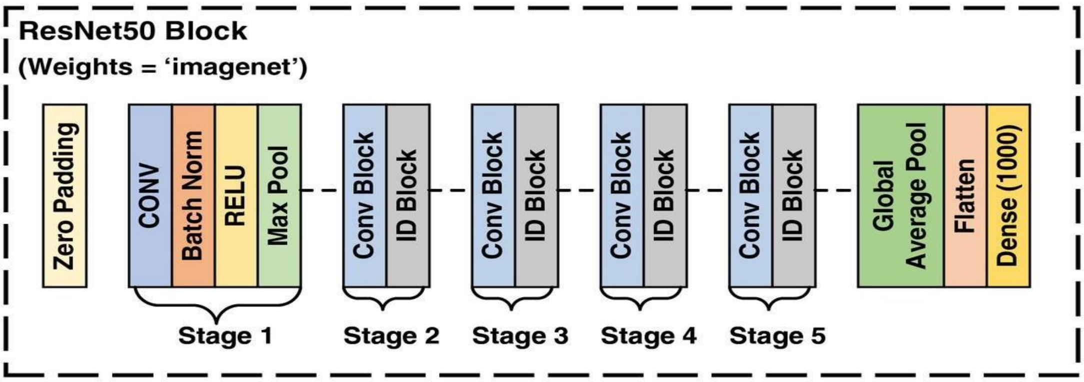

Figure 2. ResNet50 Architecture

2. **Enhanced Model Generalization through Data Augmentation**

Integrating Generative Adversarial Networks (GANs) to produce lifelike synthetic images can help overcome dataset restrictions. This enhances classification performance by improving the model's generalization to various and unknown damage circumstances.

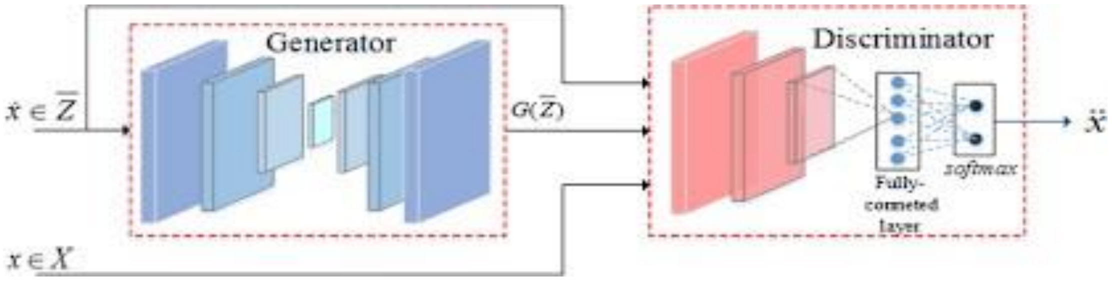

Figure 3. GAN Architecture

3. **Acceleration of Claims Processing**

Deploy a streamlined, end-to-end machine learning pipeline to significantly reduce the time required for vehicle inspections and claim evaluations, enabling faster settlements and reduced operational delays.

4. ***Reduction of Subjectivity and Human Error***

Provide standardized and objective damage assessments by minimizing reliance on human adjusters. This improves consistency across evaluations and reduces errors that can lead to overpayments or claim disputes.

5. ***Minimization of Financial Losses and Claim Leakage***

Improve the accuracy of damage classification to help insurers avoid unnecessary payouts and detect potentially fraudulent claims, thereby improving financial efficiency and reducing overall claim leakage.

6. ***Real-World Impact on Insurance Operations***

In a real-world insurance context, illustrate the usefulness of deep learning and GAN-augmented training, demonstrating how AI may result in quicker, more accurate, and more intelligent claims management decision-making.

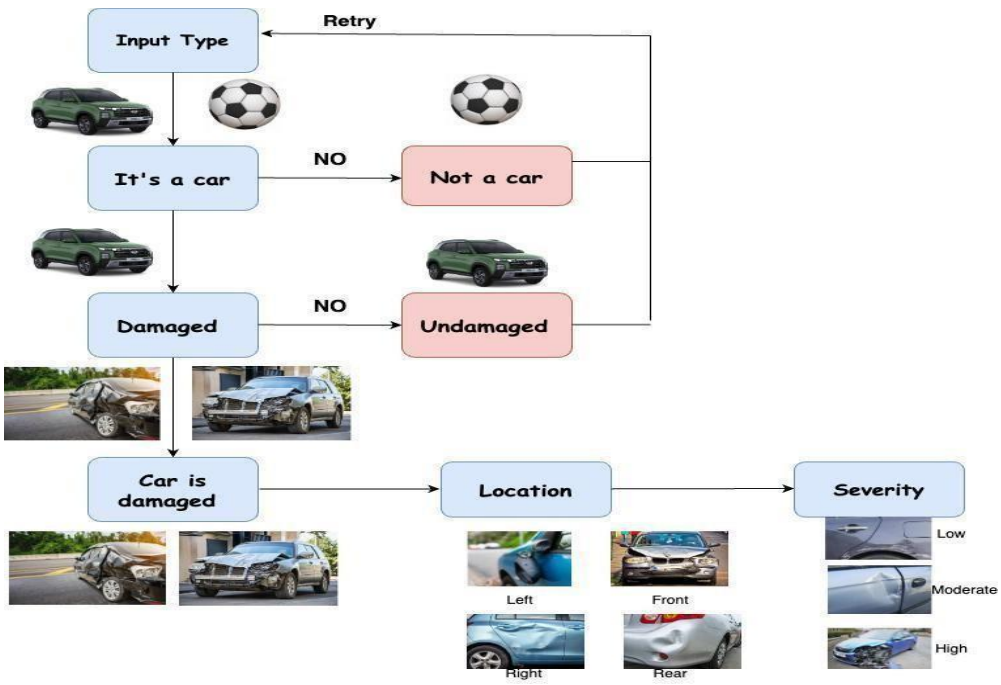

Figure 4. Workflow Representation

Methodology

This project's methodology uses a structured multi-pipeline approach to automate the assessment of vehicle damage. The procedure is broken down into multiple steps, each of which focuses on a different facet of damage assessment. The main steps in our process are listed below:

1. ***Data Collection and Preprocessing***

Initially, a dataset of vehicle photos—which may comprise both damaged and undamaged cars—must be gathered. To make sure the photos are clear and prepared for model training, this dataset has been preprocessed. To enhance model generalisation, preprocessing techniques include downsizing the pictures, normalising pixel values, and enhancing the data. To improve image quality and guarantee interoperability with deep learning models, noise is eliminated.

ImageNet Dataset:

Utilized in the transfer learning phase to pre-train ResNet50, enabling the model to take advantage of acquired characteristics from a big image dataset.

Damage Classification Dataset:

Contains labeled images of both damaged and undamaged vehicles for binary classification.

Damage Location Dataset:

Categorizes images based on where the damage is located—**front**, **rear**, or **side**—to train the model for region-specific recognition.

Damage Severity Dataset:

Includes images categorized by the level of damage severity—**minor**, **moderate**, or **severe**—used to assess the intensity of damage.

2. ***Leverage Pre-trained ResNet-50 for Feature Extraction***

Employed the basic model, ResNet-50, a deep convolutional neural network that has been previously trained on ImageNet, in addition to Deep hierarchical characteristics were extracted from car photos to form the basis for classification tasks. We froze the initial layers to enable fine-tuning while keeping general features.

3. ***Fine-Tune with L2 Regularization for Robust Learning***

To ensure the model generalizes well to unseen data, L2 regularization with a lambda value of 0.02 is incorporated during the fine-tuning phase. This technique mitigates overfitting by discouraging overly complex models that might perform well on training data but poorly on validation sets. By penalizing high-weight magnitudes, L2 regularization encourages the model to prioritize learning the most critical features relevant to vehicle damage rather than amplifying noise. This approach results in smoother and more stable training curves, which in turn improves the model's validation accuracy and overall robustness.

4. ***Use GANs for Synthetic Data Augmentation***

High-quality synthetic images of automobile damage are produced using Generative Adversarial Networks (GANs). By simulating real-world situations, these artificial samples enable the model to learn from damage types that may not have been adequately represented in the original dataset. In particular, unusual cases that are challenging to gather in big quantities, like catastrophic rear-end crashes, are added to the data with the use of GANs. The model's capacity to generalize and function well under a wider range of damage situations is improved by this artificial enlargement.

5. ***Automate Damage Presence and Severity Classification***

A binary classification model, built using features extracted from ResNet-50, is used to determine whether a car is damaged. If damage is detected, the model then extends its capability to perform multi-class classification, categorizing the severity into three distinct levels: Minor, Moderate, or Severe. Training the model on a combination of original and GAN-augmented images significantly boosts its accuracy and reliability, particularly in edge cases where conventional datasets may lack sufficient representation.

6. ***Modular Pipeline Design***

Implement a structured pipeline broken into distinct modules:
Pipe 1 – *Car Validation*: Confirms presence of a vehicle in the image.
Pipe 2 – *Damage Detection*: Identifies whether damage exists.
Pipe 3 – *Damage Localization*: Classifies damage location (Front, Rear, Side).
Pipe 4 – *Damage Severity*: Classifies the damage as Minor, Moderate, or Severe.
Each module processes its input independently but is integrated into a unified flow controlled via (pipeline.py).

7. ***Frontend Deployment with Flask***

To make the system user-accessible, a lightweight web application is developed using Flask. This frontend allows users to upload car images and receive real-time analysis through a simple interface. Upon uploading, the system displays predictions for car presence, damage detection, location of the damage, and its severity. Additionally, the interface provides user feedback through error messages when the uploaded image is unclear or does not contain a car, thus ensuring a smooth user experience and encouraging correct image input.

8. ***Reduce Claim Leakage via Accurate Prediction***

By automating the damage assessment process, the system delivers consistent and objective evaluations that insurers can trust. Unlike human assessments, which are prone to variability and oversight, this AI-driven solution minimizes errors and fraud in claims processing. As a result, it reduces claim leakage, helping insurers avoid overpayments and detect false claims. Ultimately, this leads to faster, more reliable, and cost-efficient claims handling powered by data-driven insights.

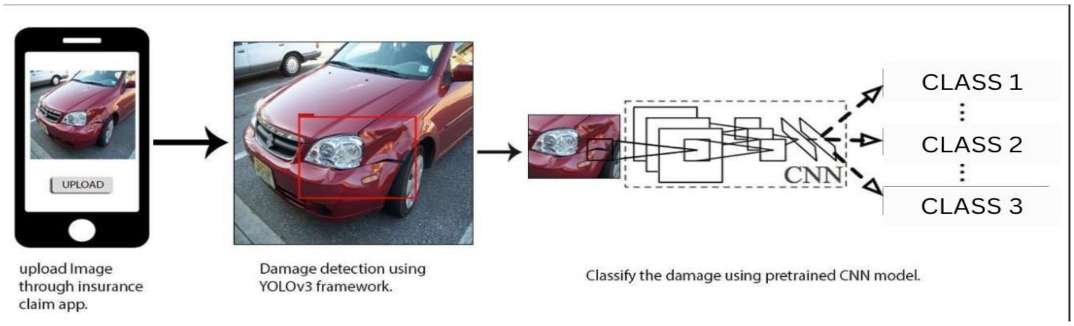

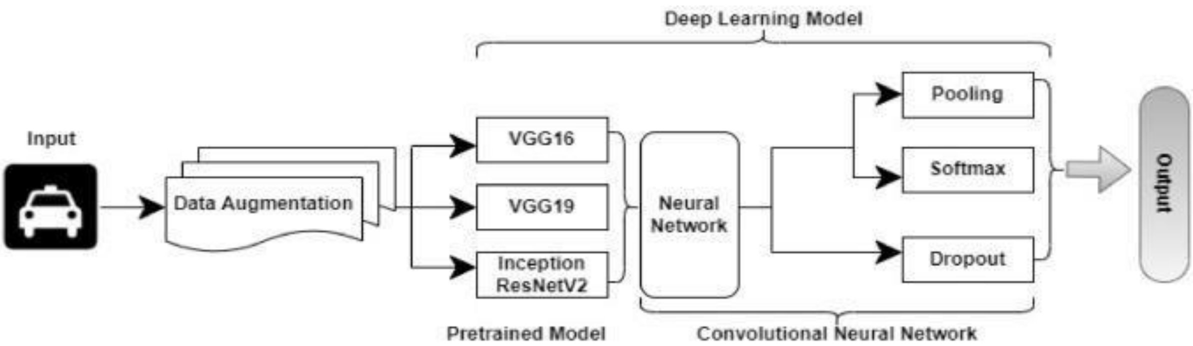

Figure 6. CNN Model Architecture Using Transfer Learning

$$\text{Softmax } \sigma(\vec{z})_i = \frac{e^{z_i}}{\sum_{j=1}^{K} e^{z_j}}$$

$$\text{Sigmoid } S(x) = \frac{1}{1 + e^{-x}}$$

Equation 1. Classifier

$$J_G = -\frac{1}{m}\Sigma_{i=1}^{m} log D(G(z_i))$$

Equation 2. GAN Generator Model

$$J_D = -\frac{1}{m}\Sigma_{i=1}^{m} log\ D(x_i) - \frac{1}{m}\Sigma_{i=1}^{m} log(1 - D(G(z_i))$$

Equation 3. GAN Discriminator Model

Chapter 4

Implementation & Results

Pipeline Implementation

Pipeline 1: Car Validation Purpose:
Ensure that the uploaded image is of a car before starting the damage assessment.

Steps:

1. **Image Preprocessing**:

The uploaded image is resized to 224x224 pixels and converted into a format suitable for deep learning models. This includes converting the image to a numpy array and normalizing it for prediction.

2. ***Prediction with resnet50:***:

The preprocessed image is fed into the resnet50, a pretrained convolutional neural network that classifies the given image into 1000 categories. The model returns a list of the top predictions along with their confidence scores.

3. ***Car Category Check***:

The model's top predictions are cross-referenced with a predefined list of car-related categories (e.g., "minivan," "sports_car," "pickup"). If any of the top predictions match a car category, the image is considered valid.

4. ***Outcome***:

- **Valid Car Image**: If the image matches a car category, the system confirms that the image is a car and proceeds to the damage assessment stage.

- **Invalid Image**: If the image does not match any car category, the system rejects it and prompts the user to upload a new image, potentially advising on better angels or lighting.

Example:

- When a car image (e.g., a **damaged car**) is uploaded, the model predicts "minivan" as one of the top categories, allowing the system to proceed.
- When a **non-car image** (e.g., a basketball player) is uploaded, the system does not find a match with the car categories and asks the user to upload a new image

	Model	**Label**	**Confidence**
0	ResNet50	minivan	0.866996
1	VGG16	minivan	0.707445
2	VGG19	sports_car	0.339943

3	VGG19	minivan	0.215945
4	VGG16	beach_wagon	0.137275
5	VGG19	beach_wagon	0.113862
6	VGG19	car_wheel	0. 083494
7	VGG16	sports_car	0.060461
8	VGG19	racer	0.049579
9	VGG16	grille	0.040196
10	ResNet50	cab	0.026963
11	ResNet50	sports_car	0.023549
12	ResNet50	grille	0.021249
13	ResNet50	car_wheel	0.020428
14	VGG16	car_wheel	0.015105

Table 2. Model comparison of VGG16, ResNet50, VGG19 on imagenet dataset

Pipeline 2: Damage Detection

Purpose:

The goal of this stage is to verify whether the input car image contains any visible damage. It acts as a binary classifier to determine the presence or absence of damage before proceeding to location and severity analysis.

Process Overview:

1. **Image Preprocessing:**

- The input image is loaded and resized to 224x224 pixels using a custom load_image() function.
- It is then normalized and preprocessed to match the input requirements of the ResNet50 model.

2. ***Prediction Using Regularized ResNet50:***

- A binary classification model based on ResNet50 architecture is used.
- The model is fine-tuned with L2 regularization to reduce overfitting and improve generalization.
- Training was performed on a dataset that combines real and GAN-generated images of damaged and undamaged cars.

3. ***Classification Logic:***

- If the predicted score is greater than 0.5, the image is classified as damaged.
- If the predicted score is 0.5 or less, the image is classified as not damaged.

4. ***System Response:***

- No Damage Detected (Score ≤ 0.5):
  The system confirms that the vehicle appears undamaged and allows the user to proceed to damage location and severity determination stages.
- Damage Detected (Score > 0.5):
  The system advises the user to verify the image and, if necessary, upload a clearer picture. Suggestions may include changing the angle, improving lighting, or zooming in for better visibility.

Example Scenarios:

- An image showing a lightly scuffed bumper receives a prediction score ≤ 0.5. The system considers it undamaged and proceeds to the next stage.
- An image of a car with a visibly crumpled front bumper receives a score > 0.5. The system requests confirmation or a better-quality image for accurate analysis.

	Precision	**Recall**	**f1 - score**	**Support**
Undamaged	0.77	0.88	0.82	230
Damaged	0.86	0.73	0.79	230
Accuracy			0.81	460
Macro avg	0.81	0.81	0.81	460
Weighted avg	0.81	0.81	0.81	460

Table 3. Classification Report of Pipe 2 (Damage Detection) Pipeline 3: Damage Location Classification

1. **Data Augmentation:**

GAN-generated images that simulate damage at different points (front, back, and side) are included to increase the model's robustness. This improves the model's capacity to generalize to uncommon or underrepresented situations while also assisting in maintaining class balance.

2. ***Model Training:***

The extracted features are passed through fully connected layers to train a multi-class classifier. The model is optimized using categorical cross-entropy loss and softmax activation to differentiate between the three damage locations.

3. ***Evaluation:***

Model performance is evaluated using confusion matrices and classification reports. Metrics such as precision, recall, and F1-score are analyzed to assess accuracy and identify any potential weaknesses in class prediction.

Example:

- An image with visible frontal damage is classified as “Front” with 92% confidence.
- Misclassifications, such as side damage labeled as rear, are analyzed for model refinement.

	Precision	**Recall**	**f1 - score**	**Support**
0	0.93	0.48	0.64	29
1	0.67	0.87	0.76	39
2	0.81	0.86	0.83	35
Accuracy			0.69	103
Macro avg	0.80	0.74	0.74	103
Weighted avg	0.79	0.76	0.75	103

Table 4. Classification Report of Pipe 3 (Damage Location)

Pipeline 4: Damage Severity Classification

Objective: Classify the severity of car damage into: Minor, Moderate, or Severe.

Steps:

1. **Feature Extraction:**

Deep visual characteristics are extracted from car photos using the ResNet50 model, which has been optimized with L2 regularization. These characteristics aid in capturing intricate patterns that are essential for identifying varying degrees of damage severity.

2. ***GAN Augmentation:***

Generative Adversarial Networks (GANs) are employed to create synthetic images that represent various levels of damage severity. This technique enhances the dataset with diverse examples and addresses class imbalance.

3. ***Model Architecture and Training:***

A lightweight classifier built with fully connected dense layers and dropout is trained on the extracted features. This setup helps the model generalize while reducing overfitting on complex visual cues of damage.

4. ***Fine-Tuning:***

To improve performance, the top layers of ResNet50 are unfrozen and fine-tuned on the severity classification task. This helps the model learn more domain-specific features related to subtle damage intensity variations.

5. ***Evaluation:***

The model's predictions are evaluated using confusion matrices and accuracy scores. These metrics help measure how effectively the model distinguishes between Minor, Moderate, and Severe damage.

6. ***Prediction Function (pipe32):***

The pipeline's prediction module takes an image and outputs one of the three severity labels. For example, a severely dented vehicle image may be confidently classified as "Severe."

Example:

- A deeply dented side panel image is predicted as "Severe" with high confidence.

	Precision	**Recall**	**f1 - score**	**Support**
0	0.77	0.83	0.80	82
1	0.49	0.59	0.53	75
2	0.83	0.64	0.72	91
Accuracy			0.70	248
Macro avg	0.70	0.68	0.68	248
Weighted avg	0.71	0.69	0.69	248

Table 5. Classification Report of Pipe 4 (Damage Severity)

Pipeline 5: Full Car Damage Assessment

Pipeline 5 integrates multiple steps to check, assess, and classify car damage in images. Here's a more detailed breakdown of how each step works.

Overview of the Pipeline Steps:

Pipe 1: Check for Car (Car Detection)

- Objective: Ensure that the image provided is of a car.
- Method: The image is resized and preprocessed, then passed through a pretrained ResNet50 model. The top predictions are compared against a list of car-related categories like sedan, minivan, or SUV.

Pipe 2: Check for Damage (Damage Detection)

- **Objective**: Determine whether the vehicle in the image has damage or not.
- **Method**: A fine-tuned ResNet50 binary classifier (trained on real and GAN-generated damaged images) returns a probability score:0.5 → Damage Detected ≤ 0.5 → No Visible Damage

Pipe 3: Location Assessment (Damage Localization)

- **Objective**: Classify the location of the detected damage.
- **Location**: The image is assessed for damage location, which is classified into **front, rear or side.**

Pipe 4: Severity Assessment (Damage Localization)

- **Objective**: Classify the severity of the detected damage.
- **Severity**: Once the location is determined, the severity is classified as **minor, moderate and severe**

Frontend Implementation (Using Flask)

The frontend for the vehicle damage assessment system is built using Flask, providing a seamless and intuitive interface for users to upload images and receive damage evaluation results. The implementation bridges advanced backend processing with a user-friendly design, ensuring accessibility while maintaining technical robustness.

Application Structure:

The application follows a modular structure:

- **pipeline.py:** Manages image processing and model predictions, handling tasks like car validation, damage detection, and classification of damage location and severity.
- **Static Files**: Stores pre-trained models and uploaded images.
- Templates (**index.html, layout.html, results.html**): HTML files that render pages dynamically based on user interactions and assessment results.
- **app.py:** The main Flask application script that connects the backend models to the frontend.

Workflow and Features:

5. **User Interaction (index.html):**
    - Users can upload an image of a damaged vehicle through a simple interface.
    - Guidance is provided to ensure images are clear and focus on major damage areas.
    - The page includes an upload button and a submission form, designed for ease of use.

6. ***Image Processing (app.py):***
    - The application validates the uploaded file format and size before saving it securely.
    - Images are passed to pipeline.py, where they undergo multiple validation steps:
        - **Car Validation:** Ensures the uploaded image depicts a car.
        - **Damage Validation:** Checks for visible damage in the image.
        - **Location and Severity Classification:** Determines where the damage is located (front, rear, side) and assesses its severity (minor, moderate, severe).

7. ***Results Display (results.html):***
    - Users receive clear feedback, including the uploaded image and results of the assessment stages.
    - Success messages indicate validation checks, damage location, and severity. If validation fails, the page offers suggestions for improving the submission (e.g., better lighting or angle).

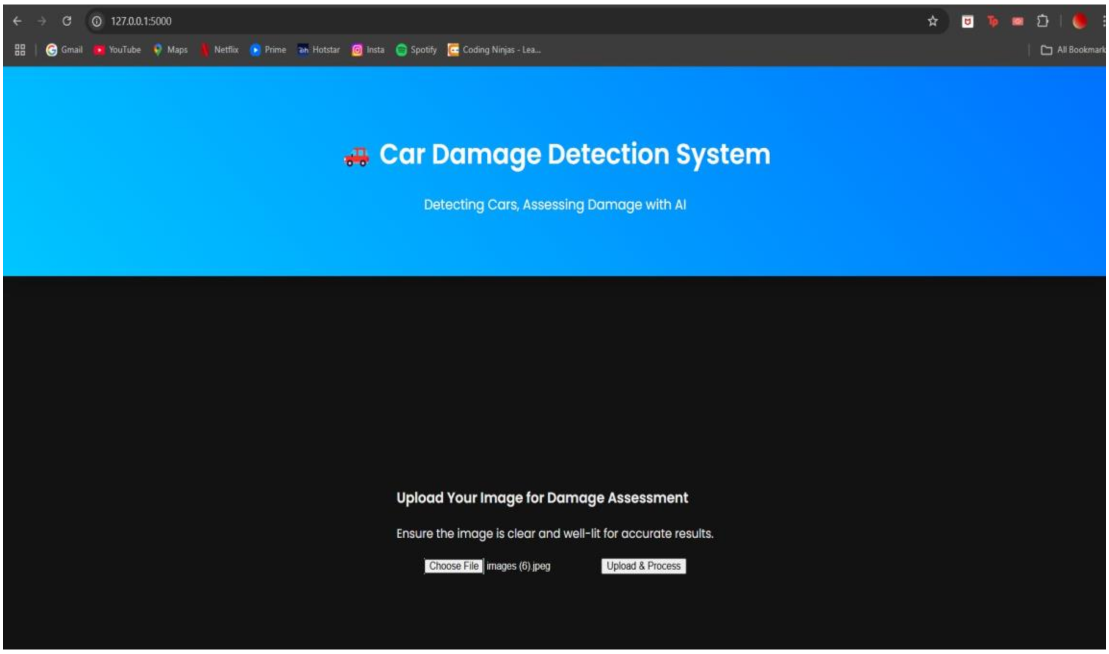

Figure 7. Website Home Page

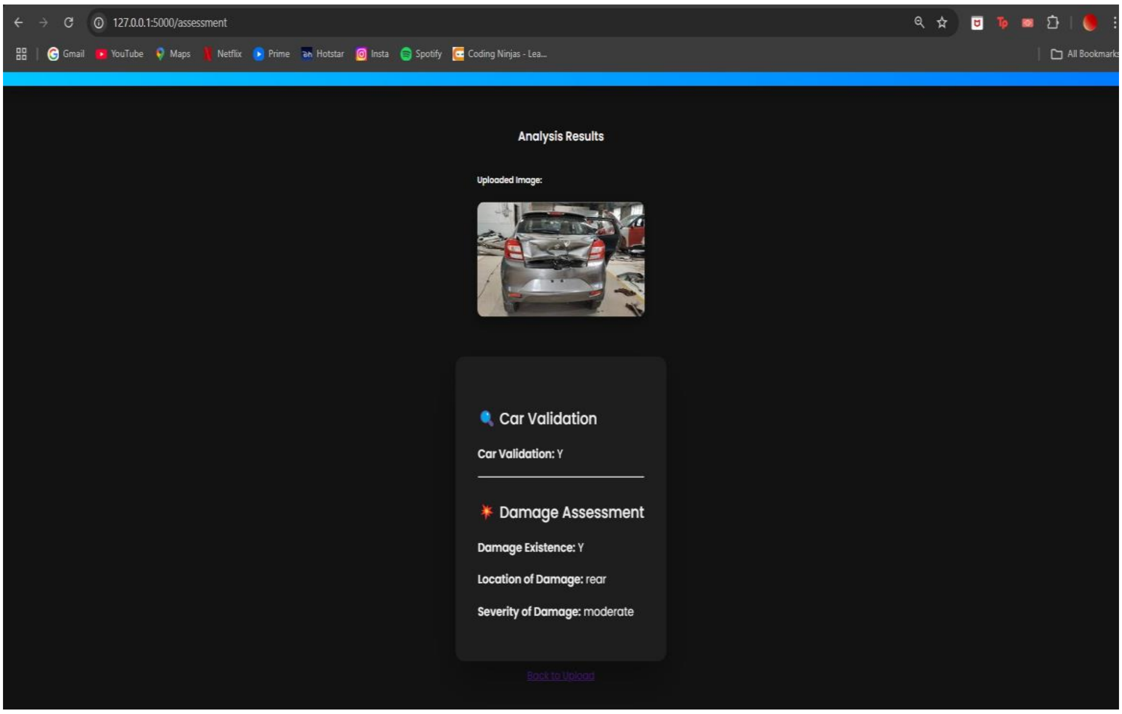

Figure 8. Website Damage Assessment Report Page

Chapter 5

RESULTS AND DISCUSSIONS

Discussion

The development of this multi-pipeline vehicle damage assessment system demonstrates how modular deep learning design can effectively cater to the nuanced demands of automated insurance claim workflows. Rather than evaluating each model solely on individual performance metrics, it is crucial to understand how the interplay between model architecture, data augmentation, regularization, and pipeline modularity contributes to a system that is not only accurate but also explainable and deployable.

1. **Vehicle Detection and Validation: Foundation of Contextual Trust**

The pipeline begins with vehicle validation, a step that might seem straightforward but is fundamental in avoiding downstream errors. A fine-tuned ResNet-50, with frozen early layers, was employed to capitalize on general image features while adapting to domain-specific cues. This significantly improved training efficiency and reduced overfitting.

Serving as a gatekeeper, this module filters out invalid or misaligned inputs, which is critical for both model performance and practical reliability. By preventing poorly framed or irrelevant images from progressing further, the system ensures damage analysis is grounded in valid data—preserving trust, especially in sensitive domains like insurance.

2. **Damage Detection: Establishing the Baseline Condition**

This binary classification module evaluates whether damage exists. Despite its simplicity, its reliability under varied lighting and perspectives is crucial. The use of synthetic data generated by GANs enriched the dataset with hard-to-collect real-world scenarios—such as light scratches or dent patterns under reflections. These variations, coupled with L2 regularization ($\lambda = 0.02$), helped reduce overfitting and encouraged broader generalization.

This stage sets the tone for subsequent modules. Without confident detection here, even a perfect localization or severity model would be rendered irrelevant.

$$L_2 \text{ regularization } = w_1^2 + w_2^2 + \ldots + w_n^2$$

Equation 4. Ridge Regression Regularization

3. **Damage Localization: Structuring the Spatial Understanding**

Once damage is detected, its location must be identified. The localization module segments damage into front, rear, and side zones, echoing how adjusters determine repair estimates. However, some confusion between rear and side views revealed the limitations imposed by fixed-perspective inputs. These findings open opportunities for improvements through multi-view datasets or spatial attention networks.

4. **Damage Severity Classification: Depth of Insight, Not Just Depth of Learning**

This module classifies damage into Minor, Moderate, or Severe—an inherently subjective task with financial implications. Class imbalance posed a major hurdle, with Severe cases underrepresented. GAN-based augmentation addressed this by creating synthetic Severe examples that balanced the training distribution and exposed the model to critical edge cases.

Regularization helped mitigate overfitting on dominant classes and facilitated feature learning beyond surface texture. The outcome was a model capable not just of identifying a dent, but of judging its impact—a distinction that insurance professionals deeply value.

5. **Modular Pipeline Design: Building Explainability and Scalability**

Opting for a modular design instead of a monolithic architecture proved advantageous. Each component was independently trainable and interpretable. If the damage severity module required re-training on new policy data, it could be done without altering the detection or localization modules.

This modularity also enabled traceability. When queried about an output, the system could walk backward—validating the presence of a vehicle, identifying the damage, and reasoning about its location and severity. This transparency is a strong asset for enterprise deployment.

6. **Real-World Integration: Feedback, Usability, and Insurance Utility**

Deployment through a Flask-based frontend allowed usability testing in realistic environments.This user-guided data improvement loop proved vital for systems operating in mobile or decentralized environments.

7. **The Role of GANs and Regularization: More Than Just Accuracy**

Both GANs and regularization were essential not just for improving raw accuracy but for achieving robust and fair learning. GANs helped fill data gaps by synthesizing rare but realistic examples, particularly for the Severe damage class. Meanwhile, regularization ensured that the models didn't simply memorize frequent patterns, allowing them to generalize to unseen vehicle types, angles, and lighting conditions.

Results

1. ***Vehicle Validation (Pipe 1: Car Detection)***

    1. **Objective**: Confirm the presence of a car in the uploaded image.
    2. **Model**: Fine-tuned ResNet-50 with frozen early layers.

3. ***Performance:***

Accuracy: 86.6%

- Observation: The model performed reliably across varied backgrounds. False negatives primarily occurred due to poor lighting or partial vehicle visibility.

2. ***Damage Detection (Pipe 2: Binary Classification)***

1. **Objective**: Determine whether the vehicle is damaged.
2. **Dataset**: Real + GAN-augmented images of damaged and undamaged vehicles.

3. ***Performance:***

Accuracy: 80.65%

4. ***Enhancements:***

- GAN augmentation improved detection of rare, subtle damages.
- L2 regularization (λ = 0.02) reduced overfitting and stabilized model training, resulting in smoother loss curves and consistent performance across epochs.

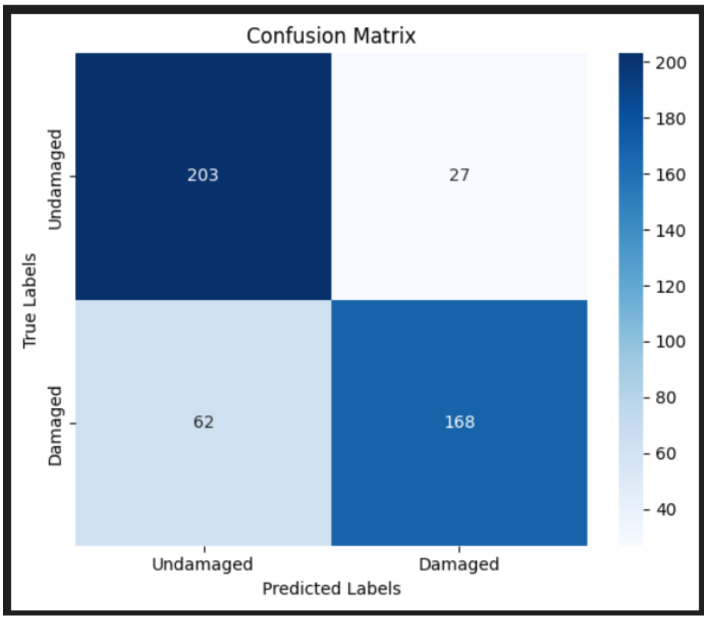

Figure 9. Confusion Matrix Of Pipe 2 (Binary Classification)

3. **Damage Localization (Pipe 3: Region Detection)**

1. **Objective**: Identify where the damage is located—Front, Rear, or Side.
2. ***Performance:***

- Accuracy: 68.9%

3. **Challenges**: Some misclassifications occurred between rear and side damages in oblique-angle photos. Region-specific training helped reduce these overlaps.

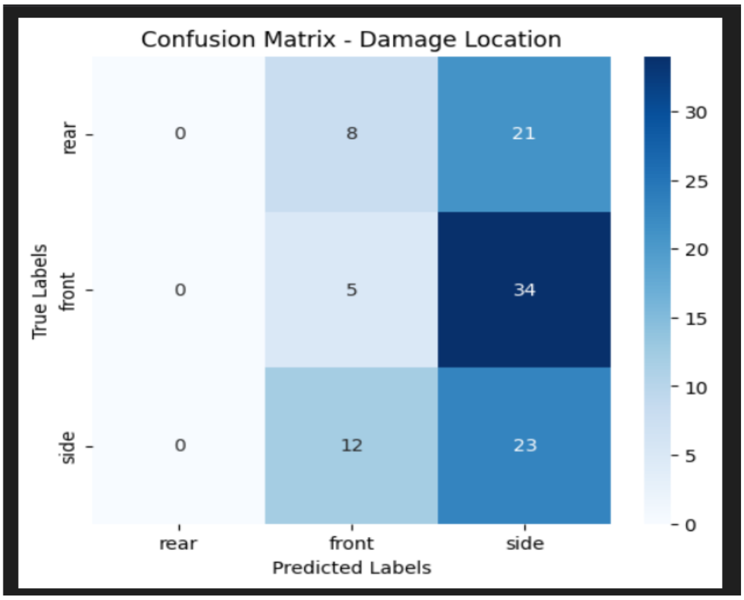

Figure 10. Confusion Matrix Of Pipe 3 (Region Detection)

4. **Damage Severity Classification (Pipe 4: Multi-Class Classification)**

1. **Objective**: Classify damage severity as Minor, Moderate, or Severe.
2. ***Performance:***
   - Accuracy: 70.16%
3. **Impact of GANs**: GAN-generated images, especially for Severe cases, corrected class imbalance and improved model sensitivity to high-intensity damage.
4. ***Regularization Outcome:***
   - L2 regularization led to higher generalization and prevented overfitting in minority classes.

Technique	Role in Training	Benefits
GANs (Synthetic Augmentation)	Generate rare and high-diversity samples (e.g. severe or unusual damages)	-Mitigated class imbalance -Enriched edge-case representation -Enhanced real-world generalisation
L2 Regularisation	Applied during fine-tuning of ResNet-50	-Suppressed overfitting -Smoothed training curves -Encouraged learning essential, low-complexity features

Table 6. Advantages of GANs and L2 Regularization

Future Works

To further enhance the capabilities and real-world applicability of the car damage assessment system, several promising directions are proposed. First, incorporating multi-view imaging will significantly improve damage detection accuracy by allowing the system to analyze vehicle damage from multiple angles, thus reducing blind spots and providing a more comprehensive assessment. Additionally, experimenting with advanced deep learning architectures such as Vision Transformers (ViT) could lead to better performance, especially in identifying complex or subtle damage patterns that traditional convolutional networks might miss.

Expanding the dataset is another crucial step. By collecting a larger and more diverse set of real-world vehicle images-covering various makes, models, lighting conditions, and damage types-the model's robustness and generalizability can be greatly improved. Making the system accessible through mobile and web platforms is also a priority, as this would enable users, insurance agents, and repair shops to easily upload images and receive instant assessments, streamlining the claims and repair process.

Furthermore, integrating repair cost estimation into the pipeline would provide users with not only the location and severity of damage but also an approximate repair cost, making the evaluation more actionable and informative. Finally, as data privacy and security are paramount, especially in insurance applications, exploring the use of blockchain technology could ensure that claims data remains secure, tamper-proof, and transparent, building greater trust among users and stakeholders. These future enhancements aim to make the system more accurate, user-friendly, and reliable for real-world deployment.

REFERENCES

1. Gandhi, Ritik. "Deep learning based car damage detection, classification and severity." International Journal 10.5 (2021).
2. Kadakia, Sarang, Om Chavan, Prem Doshi, and Prachi Tawde. "Dent Detection And Price Prediction For Cars: A Comprehensive Analysis." International Journal of Creative Research Thoughts, vol. 11, no. 10, 2023.
3. Kyu, Phyu Mar, and Kuntpong Woraratpanya. "Car damage detection and classification." Proceedings of the 11th international conference on advances in information technology. 2020.
4. Liu, Yucheng, Jun Zhang, and Song Li. "Automated Vehicle Damage Inspection Using Deep Learning and Multi-View Imaging." IEEE Transactions on Intelligent Transportation Systems, vol. 23, no. 6, 2022, pp. 5678–5690.
5. Mahyoub, Mohamed, Panos Liatsis, Friska Natalia, Sud Sudirman, and Abdulmajeed Hammadi Jasim Al-Jumaily. "Data Augmentation Using Generative Adversarial Networks to Reduce Data Imbalance with Application in Car Damage Detection." 2023 15th International Conference on Developments in eSystems Engineering (DeSE), pp. 480-485, 2023.
6. Patil, Kalpesh R., and Masaaki Iiyama. "Deep learning models to predict sea surface temperature in Tohoku region." IEEE Access 10 (2022): 40410-40418.
7. Wang, Xiu, and Ankush Gupta. "Enhancing Automotive Damage Detection with Vision Transformers." Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 2023, pp. 123–130.
8. Warungse, Priti, Atharva Kasar, Apurva Kshirsagar, Geeta Hade, and Nishant Khandhar. "Car Damage Detection Using Computer Vision." Journal of Emerging Technologies and Innovative Research, vol. 10, no. 11, 2023.