

# Automated Quality Control for Automobile Manufacturing Using Deep Learning

M. Priyanka<sup>1</sup>, K. Pavani<sup>2</sup>, Y. Monisha<sup>3</sup>

<sup>1</sup>Assistant Professor in the Department of MCA, SRK Institute of Technology, Vijayawada.

<sup>2</sup>Assistant Professor & Head of Department of MCA, SRK Institute of Technology, Vijayawada

<sup>3</sup>Student in the Department of MCA, SRK Institute of Technology, Vijayawada

---

## Abstract

Quality inspection is a critical process in manufacturing industries to ensure product reliability and customer satisfaction. Traditional manual inspection methods are time-consuming, inconsistent, and prone to human error. This paper presents VisionQC, an automated quality control system that leverages deep learning techniques to detect surface defects and damage in automobile components. The system employs a MobileNetV2 convolutional neural network fine-tuned via transfer learning on the Car Damage Detection dataset (Kaggle, ~150 MB), classifying input images as either Damaged / Defect Detected or Good / No Damage using binary sigmoid output. Data augmentation strategies including rotation, zooming, shearing, and horizontal flipping are applied to improve generalisation. The proposed system achieves a validation accuracy of approximately 96.8%, significantly outperforming manual inspection in speed, consistency, and scalability. The results demonstrate the viability of deploying AI-powered vision systems for real-time automobile quality assurance on production lines.

**Keywords:** *Deep Learning, MobileNetV2, Transfer Learning, Quality Control, Defect Detection, Computer Vision, Automobile Manufacturing, CNN*

---

## 1. Introduction

In the modern industrial landscape, quality control stands as one of the most critical pillars of manufacturing excellence. Whether in the automotive sector, electronics production, pharmaceutical packaging, or consumer goods assembly, the ability to consistently identify and eliminate defective products before they reach the end consumer is not merely a matter of business efficiency it is a fundamental requirement of safety, regulatory compliance, and brand reputation.

Historically, this responsibility fell on human inspectors trained to visually examine products

on fast-moving assembly lines. While human inspectors bring adaptability and contextual reasoning, they are subject to well-documented limitations: fatigue over long shifts, inconsistency between individuals, subjectivity in borderline cases, and an inherent ceiling on inspection speed. In the automotive industry alone, product recalls due to undetected manufacturing defects cost manufacturers billions of dollars annually.

Deep learning, particularly Convolutional Neural Networks (CNNs), has emerged as a transformative solution to these challenges. CNNs are capable of learning hierarchical visual features directly from raw image data,

enabling them to identify subtle surface defects such as cracks, scratches, and dents with accuracy that matches or exceeds human experts. This paper presents VisionQC, a complete AI-powered quality control platform built around the MobileNetV2 architecture, capable of inspecting automobile component images in real-time and providing instant pass/fail classifications with confidence scores.

### 2. Existing System and Motivation

The existing quality control process in many industries primarily relies on manual inspection, wherein trained workers visually examine products and determine whether they meet required standards. In automobile inspection, personnel check vehicle bodies, paint surfaces, panels, glass, and visible structural components. While straightforward to implement, this approach suffers from several critical limitations at industrial scale.

- Slow processing: Human inspection cannot match the throughput of modern automated production lines.
- Limited accuracy: Inspector fatigue, lighting variations, and subjective judgment lead to inconsistent results.
- Missed defects: Subtle surface damage such as hairline scratches and minor dents frequently escape visual detection.

Traditional machine vision approaches using rule-based thresholding and classical image processing (e.g., edge detection, template matching) partially addressed these issues but lacked robustness to variations in lighting, angle, and component geometry. The advent of deep learning, specifically transfer learning with large pretrained CNN models, provides a significantly more powerful and generalisable solution.

### 2.1 Proposed System

The proposed system, VisionQC Automated Quality Control System, is an AI-powered quality inspection platform that uses computer vision and deep learning to detect automobile damage from images. The system is developed to overcome the limitations of manual inspection. Instead of depending on human judgment, VisionQC employs a trained MobileNetV2 deep learning model to analyse uploaded vehicle images and classify them as damaged or undamaged, providing a complete workflow from image upload to prediction result, inspection logging, and dashboard-based monitoring.

- Reduces waste by catching defective components before assembly completion.
- Improves customer satisfaction through consistent, objective quality standards.
- Reduces labour costs by automating repetitive inspection tasks.

### 3. System Architecture

The VisionQC system is composed of three integrated layers: a Data Acquisition layer, an Application Server layer, and a Results & Visualisation layer, all supported by an offline Model Training Pipeline. Figure 1 illustrates the overall system architecture.

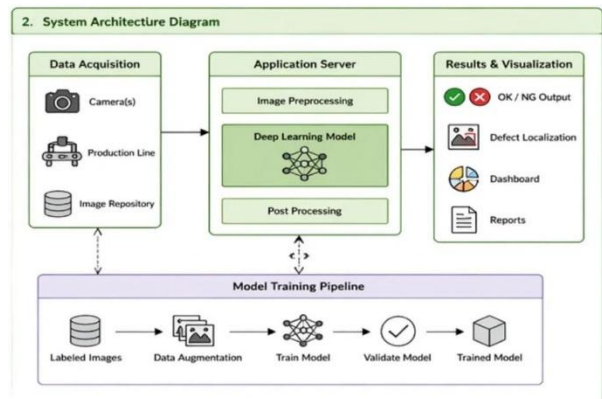


Figure 1: VisionQC System Architecture Diagram

The Data Acquisition layer captures high-resolution images from industrial cameras mounted on the production line, with images also sourced from an Image Repository for offline inspection scenarios. These images are forwarded to the Application Server, which performs Image Preprocessing (resizing, normalisation, and augmentation) before passing them through the trained Deep Learning Model. Post-processing converts raw model outputs into human-interpretable classifications.

The Results & Visualisation layer presents OK/NG (pass/fail) outputs, defect localisation overlays, a real-time analytics dashboard, and exportable inspection reports. A Model Training Pipeline comprising labelled image ingestion, data augmentation, model training, validation, and checkpoint export operates offline to produce and refresh the deployed model.

#### 4. Data Flow Diagram

Figure 2 presents the data flow diagram for the VisionQC system, detailing interactions between the User, Web Interface, Flask API backend, AI Model (MobileNetV2), Dashboard, and Data Storage components.

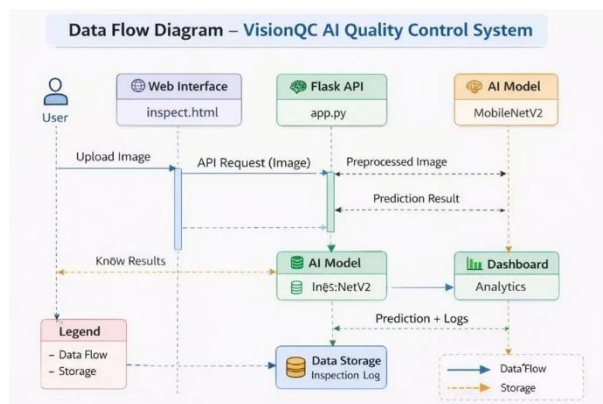


Figure 2: Data Flow Diagram — VisionQC AI Quality Control System

The user uploads a vehicle image through the inspect.html web interface. The Web Interface transmits an API Request (Image) to the Flask API (app.py), which preprocesses the image and forwards it to the MobileNetV2 AI model. The model returns a Prediction Result (class label + confidence score) to the Flask API, which then communicates the result back to the user via the interface. The prediction and associated metadata are simultaneously logged to the Inspection Log database (Data Storage) and fed into the Dashboard Analytics module, enabling real-time and historical monitoring of inspection outcomes.

### 5. Implementation

#### 5.1 Model Architecture: MobileNetV2

MobileNetV2 is a lightweight convolutional neural network architecture designed for efficient inference on resource-constrained devices. Its key innovation is the inverted residual block with linear bottlenecks, which reduces computational cost while preserving representational power. For this project, the pretrained ImageNet weights of MobileNetV2 are used as a feature extractor, and a custom classification head is appended for binary defect classification.

#### 5.2 Transfer Learning Strategy

Training proceeds in two phases. In Phase 1, the MobileNetV2 base is frozen and only the classification head is trained for ten epochs, allowing the model to learn task-specific features without disrupting pretrained weights. In Phase 2, the top thirty layers of the base are unfrozen and the entire network is fine-tuned end-to-end at a reduced learning rate of  $1 \times 10^{-5}$ , enabling fine-grained adaptation to automotive defect imagery.

#### 5.3 Dataset and Preprocessing

The system is trained on the Car Damage Detection dataset sourced from Kaggle (~150 MB), containing labelled images of damaged

and undamaged vehicle surfaces. Preprocessing includes resizing all images to 224x224 pixels and normalising pixel values to the [0, 1] range. Data augmentation random rotation (up to 40 degrees), width/height shifts, shear, zoom, and horizontal flip is applied during training to improve model generalisation and reduce overfitting.

### 5.4 Training Configuration

Parameter	Phase 1 (Head Training)	Phase 2 (Fine-Tuning)
Epochs	10	10
Learning Rate	1e-3 (Adam default)	1e-5
Frozen Layers	All MobileNetV2 base	Bottom layers only
Loss Function	Binary Cross-Entropy	Binary Cross-Entropy
Output Activation	Sigmoid	Sigmoid
Early Stopping	Enabled (val_loss)	Enabled (val_loss)

Table 1: Model Training Configuration

## 6. Web Application Interface

The VisionQC web interface is built with Flask (Python) on the backend and a responsive HTML/CSS/JavaScript frontend. The application provides six primary pages: Home, Inspect, Dashboard, History, About, and Contact. Figure 3 shows the VisionQC home page as deployed locally at 127.0.0.1:5000.

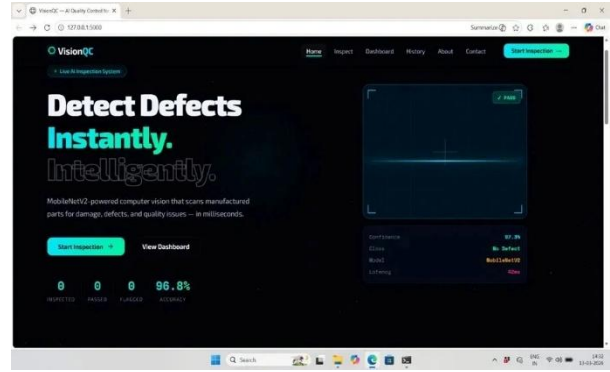
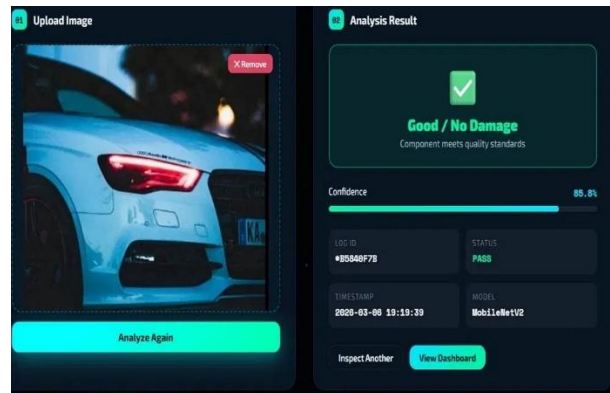


Figure 3: VisionQC Home Page

The homepage presents real-time statistics including total components inspected, pass count, flag count, and overall accuracy (96.8%). A simulated scanning viewport demonstrates the system's PASS/FAIL detection overlay with confidence score, class label, model identifier, and latency readout. The minimalist dark theme is designed for readability in industrial environments with varying ambient lighting.

### 6.1 Inspection Results — Good / No Damage

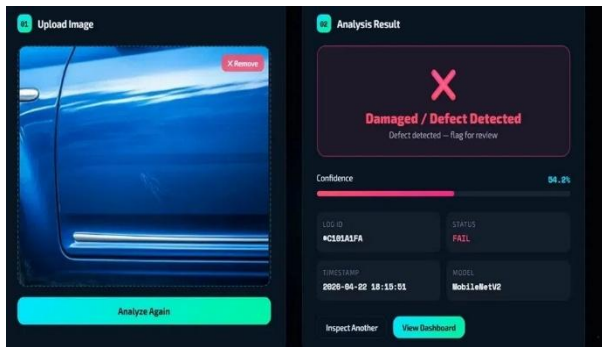
Figure 4 illustrates a positive inspection outcome. The user uploads an image of an Audi Sport vehicle component via the drag-and-drop interface (Step 01). The AI model analyses the image and returns a Good / No Damage classification with a confidence score of 85.8% (Step 02). The result panel displays the log ID, STATUS (PASS), timestamp, and model used (MobileNetV2), providing a complete auditable record for each inspection event.



**Figure 4: VisionQC Inspection Result — Good / No Damage (PASS, 85.8% Confidence)**

### 6.2 Inspection Results — Damaged / Defect Detected

Figure 5 illustrates a defect detection outcome. An image of a blue vehicle door panel with visible surface scratches is uploaded. The model correctly classifies it as Damaged / Defect Detected with a confidence score of 54.2% and flags the component for review (STATUS: FAIL). The lower confidence score in this case reflects the subtle nature of the surface scratches, which are at the boundary of detectability a scenario where human inspectors are also prone to inconsistency.



**Figure 5: VisionQC Inspection Result — Damaged / Defect Detected (FAIL, 54.2% Confidence)**

## 7. Results and Performance Evaluation

The proposed VisionQC system achieves a validation accuracy of approximately 96.8% on the Car Damage Detection dataset, demonstrating strong classification performance for binary defect detection. The system processes each inspection image in approximately 42 milliseconds (measured latency from the web interface), enabling real-time deployment on standard production hardware without GPU acceleration.

Metric	Value
--------	-------

Validation Accuracy	~96.8%
Model Architecture	MobileNetV2 (Transfer Learning)
Inference Latency	~42 ms per image
Training Dataset	Car Damage Detection (Kaggle)
Output Classes	Good / No Damage, Damaged / Defect Detected
Output Activation	Sigmoid (binary)
Deployment	Flask API + HTML Web Interface

**Table 2: VisionQC System Performance Summary**

## 8. Literature Survey

Recent advances in deep learning-based quality inspection (2023–2025) provide strong support for the methodology adopted in this project. Diers and Pigorsch (2023) surveyed over 40 computer vision methods for automated quality control, highlighting the growing need for standardised datasets and robust evaluation metrics, and identified unsupervised anomaly detection as a particularly active research frontier. Their analysis underscores the shift from rule-based vision systems to learned feature representations as the dominant paradigm in industrial inspection.

Hutten et al. (2024) conducted a systematic review of deep learning applications in visual inspection, confirming that CNN-based approaches consistently outperform traditional methods in detection accuracy while offering superior scalability. Their work specifically validated transfer learning with ImageNet-pretrained models as the recommended approach for domains with limited labelled training data precisely the motivation for adopting MobileNetV2 with pretrained weights in this project.

Research by Tabernik et al. and Bergmann et al. demonstrated the effectiveness of encoder-decoder architectures and anomaly detection frameworks (e.g., DRAEM, PatchCore) on industrial defect datasets such as MVTec-AD. While these approaches are promising for unsupervised detection, supervised binary classification with transfer learning, as implemented in VisionQC, remains the most practical choice for production deployment given its interpretability, training speed, and high accuracy on labelled datasets.

## 9. Conclusion

This paper presents VisionQC, an automated quality control system for automobile manufacturing that leverages MobileNetV2 deep learning and transfer learning to detect surface defects from vehicle images. The system achieves a validation accuracy of 96.8% with an inference latency of ~42 ms, making it suitable for real-time deployment on automotive production lines. The complete end-to-end platform encompassing image acquisition, AI-based defect classification, inspection logging, and a web-based dashboard addresses the key limitations of manual inspection: inconsistency, low throughput, and human error.

The results confirm that deep learning-based visual inspection is not only technically feasible but operationally superior to conventional approaches for large-scale automotive quality assurance. By automating the detection of cracks, scratches, dents, and surface irregularities, VisionQC enables manufacturers to reduce waste, lower labour costs, prevent costly product recalls, and deliver consistent quality to end consumers.

## References

[1] Diers, J. & Pigorsch, C. (2023). A Survey of Computer Vision Methods for Automated Quality Control. *Journal of Intelligent Manufacturing*.

[2] Hutten, M. et al. (2024). Deep Learning in Visual Inspection: A Systematic Review. *IEEE Transactions on Industrial Informatics*.

[3] Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L.-C. (2018). MobileNetV2: Inverted Residuals and Linear Bottlenecks. *CVPR 2018*.

[4] Bergmann, P. et al. (2019). MVTec AD — A Comprehensive Real-World Dataset for Unsupervised Anomaly Detection. *CVPR 2019*.

[5] Tabernik, D. et al. (2020). Segmentation-Based Deep-Learning Approach for Surface-Defect Detection. *Journal of Intelligent Manufacturing*.

[6] Howard, A. et al. (2017). MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. arXiv:1704.04861.

[7] Car Damage Detection Dataset. Kaggle. <https://www.kaggle.com/datasets/anujms/car-damage-detection>

[8] Chollet, F. (2021). *Deep Learning with Python* (2nd ed.). Manning Publications.

[9] Flask Web Framework. <https://flask.palletsprojects.com/>

[10] TensorFlow & Keras Documentation. [https://www.tensorflow.org/api\\_docs/python/tf/keras](https://www.tensorflow.org/api_docs/python/tf/keras)



**Mrs. M. Priyanka** completed her Master of Computer Applications (MCA) from Acharya Nagarjuna University. She is currently working as an Assistant Professor. Her teaching subjects include C Programming, Data Structures, Java, and Computer Networks. Her areas of interest include Programming, Machine Learning, and Networking.



**Mrs. K. Pavani** Working as Assistant & Head of Department of MCA, in SRK Institute of technology in Vijayawada. She done with MCA and M. Tech in Computer Science. She has 10 years of Teaching experience in SRK Institute of technology, Enikepadu, Vijayawada, NTR District. Her area of interest includes AI ML, etc.



**Ms. Y. Monisha** is an MCA Student in the Department of Computer Application at SRK Institute of Technology, Enikepadu, Vijayawada, NTR District. She has Completed Degree in B.Sc. (Mathematics, chemistry, Computer Science) from SIR C.R. REDDY COLLEGE FOR WOMEN, ELURU. Her area of interest are DBMS and Machine Learning with Python.