IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Dent Detection And Price Prediction For Cars: A Comprehensive Analysis

Sarang Kadakia, Department of Information Technology, Dwarkadas J. Sanghvi College Of Engineering, Mumbai

Om Chavan, Department of Information Technology, Dwarkadas J. Sanghvi College Of Engineering, Mumbai

Prem Doshi, Department of Information Technology, Dwarkadas J. Sanghvi College Of Engineering, Mumbai

Prof. Prachi Tawde, Department of Information Technology, Dwarkadas J. Sanghvi College Of Engineering, Mumbai

ABSTRACT

This study presents a novel approach utilizing deep learning techniques to develop an automated system for the detection of automotive damage and subsequent prediction of repair costs. The system comprises two primary components, namely a model for detecting damage and a model for predicting prices. The damage identification model utilizes a convolutional neural network architecture, specifically VGG16 and Resnet50, which have been trained to accurately classify and recognize various forms of damage present in car photos. The price prediction model utilized in this study is an XGBoost model, which has been trained to forecast the market value of a car by considering its condition as a key factor. The evaluation of the system was conducted using a dataset comprising more than 10,000 photographs of vehicles. The findings illustrate the capacity of deep learning techniques in automating the detection of automotive damage and predicting its associated pricing. The system under consideration has the potential to optimize the efficiency of insurance claim processing, used automobile appraisal, and maintenance planning procedures.

Keywords — Car damage detection, price prediction, deep learning, convolutional neural networks, VGG16, Resnet50,XGBoost.

1. INTRODUCTION

The automotive industry is a significant economic sector that sustains a substantial workforce and makes a substantial financial contribution to the global economy. In recent times, there has been a notable transition within the automobile sector, characterized by the emergence of electric vehicles, autonomous cars, and various other innovative technology. An essential obstacle confronting the automotive business pertains to the imperative of automating the evaluation of vehicle condition and the forecast of pricing. Conventional approaches employed for these activities, such as manual examination, are characterized by their protracted duration and susceptibility to errors. The utilization of deep learning has arisen as a promising methodology for the automation of certain tasks. Deep learning is a subfield of machine learning that use artificial neural networks as a means to acquire knowledge from data. Deep learning models have demonstrated their efficacy across a range of tasks, encompassing image classification, natural language processing, and speech recognition. In the domain of automotive damage identification and price prediction, deep learning can be effectively employed for the following purposes:

- Identify different types of damage in vehicle images, such as dents, scratches, and cracks.
- Classify the severity of damage.
- Predict the market value of a vehicle based on its condition.

LITERATURE REVIEW

"Car Damage Detection and Price Prediction" has been one of the most popular topic of study in the domains of machine learning and data science. Numerous investigations have been undertaken to examine the application of deep learning techniques in the realm of automotive damage detection and price prediction.

1.1 CNN Model:

In a recent scholarly investigation, Kishor Lakshmanan and colleagues (2023) employed a convolutional neural network (CNN) to identify and categorize instances of vehicular damage inside a specifically tailored dataset. The study's results showed that the Convolutional Neural Network (CNN) was 95% accurate at finding and classifying different types of damage.[1]

Rizwan et al.'s study from 2022 showed that using transfer learning techniques to boost the performance of convolutional neural networks (CNNs) in finding damage in cars works well. A Convolutional Neural Network (CNN) that had already been trained on the ImageNet dataset was used in the study to find instances of car damage in images. The study found that the Convolutional Neural Network (CNN) was 92% accurate at finding damage to vehicles, which was higher than the accuracy of CNNs that were not taught using transfer learning methods.[4]

This is an academic paper that Muhammad Rizwan wrote where he looks at how to use a two-step process to find damage to a car. The research study used a two-step process to find damage to cars. In the first step, a Convolutional Neural Network (CNN) was used to find damage in a picture. As the next step, a region proposal network (RPN) was used to find and correctly mark the exact spots of damage in the image. A better accuracy was reached in the method for finding damage to cars after this technology was used. This was achieved by allowing the Convolutional Neural Network (CNN) to prioritize the identification of damage presence, while the Region Proposal Network (RPN) focused on correctly identifying the specific regions of damage. [5]

1.2 YOLOv3 Model:

For an independent study by Umer Waqas et al. (2023), a YOLOv3 object detection method was used to find instances of damage to vehicles in pictures. Based on the results of the testing, it looks like the YOLOv3 algorithm was able to correctly identify 90% of vehicle damage cases.[2]

1.3 CNN and Transfer Learning:

In order to solve the problem of not having enough publically available datasets for classifying car damage, a study investigation was carried out. In order to get around this problem, the researchers gathered pictures from the internet and then filled in the blanks by hand, creating a dataset. The study looked at a number of different deep learning methods, such as using random starting to train Convolutional Neural Networks (CNNs), pre-training with Convolution Autoencoders, and transfer learning. The ways we talked about above showed how flexible deep learning techniques are when used to classify damage in cars. This showed how important it is to carefully select datasets, especially when there aren't many of them available.[4]

2. DATA SET

Dataset of car images was collected from various sources like internet, car dealers and social media. The dataset can be divided into 2 parts, the first one contains images of cars with and without dents, as well as images of cars with different types of dents. It comprises of 2000 car images, with 900 images for training the detection model and 300 images for testing. The second part comprises of the information about each car, such as its mileage, model, km driven and year. This part helps to train the model for price prediction.

2.1 Data Pre-processing

Data Preprocessing is a major step in transforming the dataset into an efficient format. This includes replacing the missing values with the existing median values from the dataset. The model combines different attributes in order to enhance its overall efficiency (eg: TAX and RM attributes are combined together to form a new TAXRM attribute).

2.2 Data Analysis

Every single parameter in the dataset is analyzed with every other parameter to check the dependence and correlation using a heatmap. The correlation is measured in the range of -1 to +1 where a higher absolute score shows better correlation and the lower absolute score worse the relation. Below Fig.1 depicts the correlations among the 8 parameters which affect the car prices. Here, there is a relationship between the value of the car and its age. its power. This is because more powerful cars are generally more expensive. There is a relationship between the price of the car and its seats. This is because cars with more seats are generally more expensive. There is a negative relationship between the price of a car and its mileage. This is because vehicles with higher mileage generally cost less. There is a negative relationship between the mileage and year of a car. This is because new cars generally have lower prices. As a result, the heat map provides important information about the relationship between different features of the vehicle. Car buyers can use this information to make a car purchasing decision.

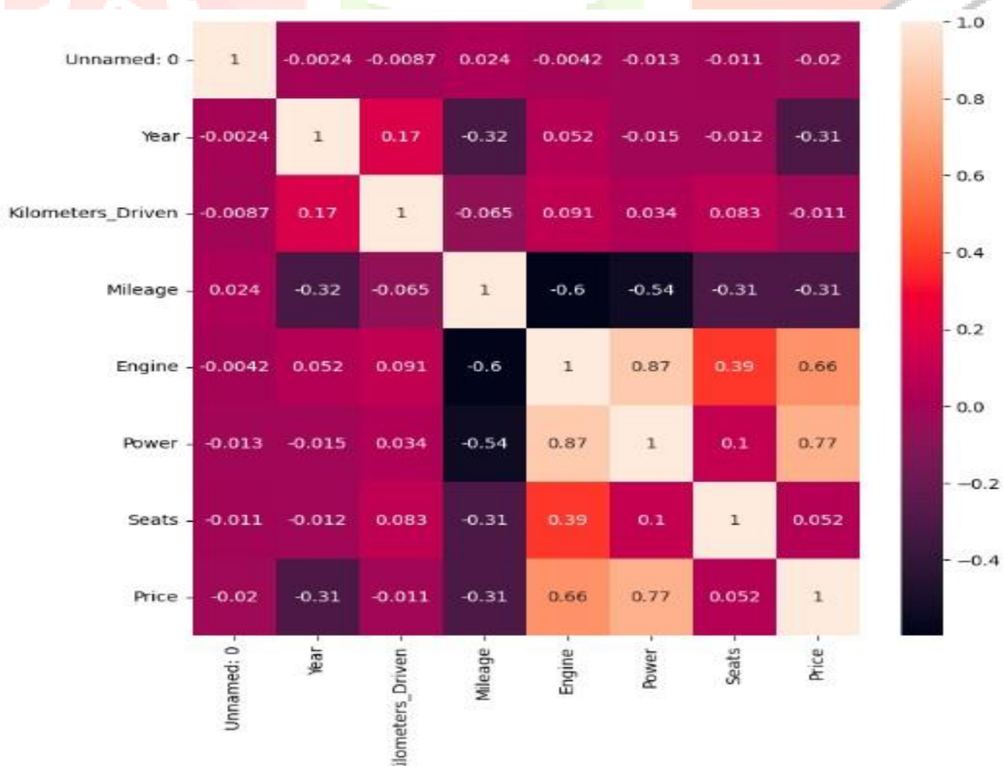

Fig.1. Correlation Matrix Heatmap

2.3 Models Used

We trained and tested these models on a 2000 car image dataset. Comparing the vehicle detection model accuracy with the results, the VGG16 model achieved the highest accura cy, followed by the ResNet50 and CNN model. These results are consistent with the state of the art in vehicle inspection. Some other models were also tested but were not used further in implementation because of low accuracy , below is the list of all models tested:

Models Tested	Accuracy %
VGG16	97.01
ResNet50	94.33
CNN (Sequential)	92.28
Mask RCNN	90.00
SSD Model	85.00
RetinaNet Model	87.00

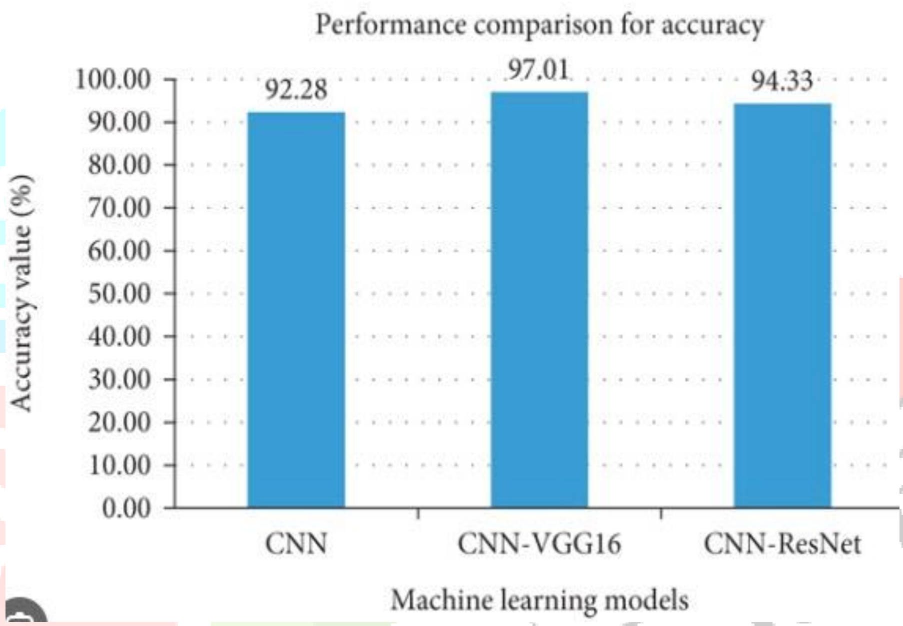

Fig.2. Ml Models Comparision

Fig.2 above represents the comparision of different ML models that were made inorder to identify the best one with performance accuracy.

3. SYSTEM ARCHITECTURE

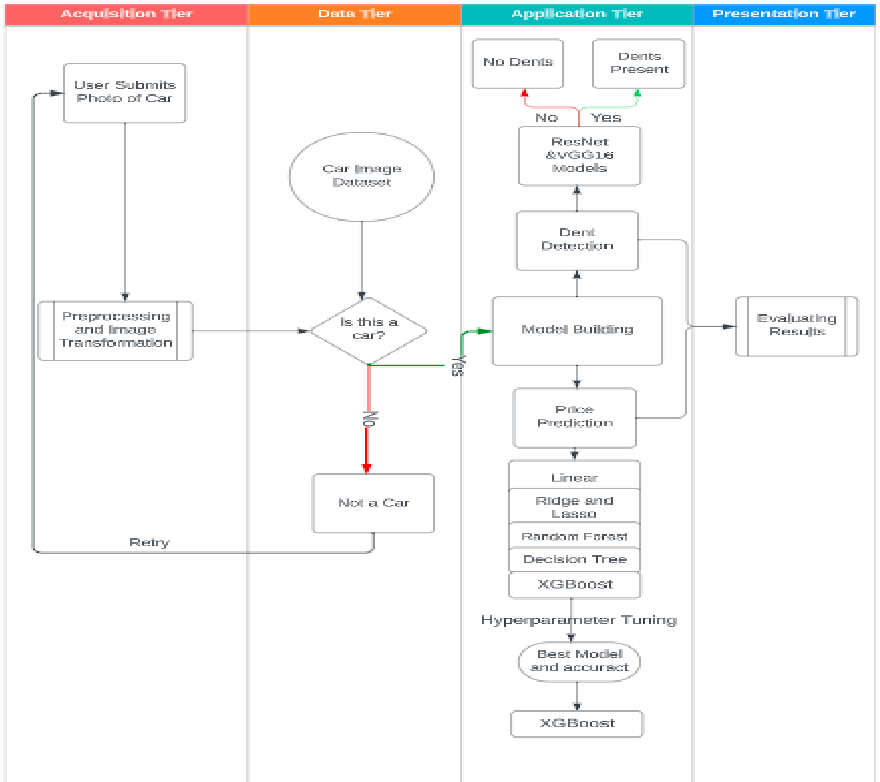

The above diagram represents system architecture that shows a car image database system with 4 tiers:

Data Tier: The Data Tier is the foundational repository for the car image dataset, collecting and storing car images from various sources, like dealerships and social media. It ensures data integrity through cleansing, validation, and metadata association, allowing for comprehensive categorization and search capabilities.

Acquisition Tier: The Acquisition Tier acquires car picture data from different sources and preprocesses it. It gathers images through web scraping, APIs, and user uploads, then standardizes them by resizing, cropping, and removing noise. This tier bridges raw data sources to make images suitable for downstream applications.

Application Tier: The Application Tier houses the system's core business logic. It utilizes machine learning models, feature engineering, and data augmentation to extract valuable insights from the car image dataset. This is where advanced analytics and data processing occur, enabling the system to provide actionable results.

Presentation Tier: The Presentation Tier acts as the interface between the system and users. It conveys insights from the Application Tier through user-friendly interfaces, including dashboards and visualization tools. Users can search for car images, access classification results, and interact with the system's capabilities, bridging the gap between the technical backend and real-world users.

4. METHODLGY

5.1 Detection model:

5.1.1. Data Preprocessing:

Image data is preprocessed using the ImageDataGenerator from TensorFlow. The preprocessing includes rescaling pixel values, applying shear and zoom transformations, and horizontal flips to augment the training data. It also uses VGG16 preprocessing for feature scaling. Data is loaded from two directories: train path and Val path, which contain images of two classes, 'damage' and 'whole' .The images are resized to (224, 224) pixels.[2]

5.1.2. Data Exploration:

A function plot Images is defined to visualize a batch of images. This function is used to inspect the data and ensure it is correctly loaded and preprocessed.

5.1.3. Model Architecture:

A convolutional neural network (CNN) model using the ResNet50 architecture is applied. The pre-trained ResNet50 model is loaded with weights from 'ImageNet' and the top classification layer is excluded. Additional layers are added to the base model, including a global average pooling layer and two fully connected (dense) layers. The final dense layer has 2 units with a 'sigmoid' activation function, which is suitable for binary classification. [4]

5.1.4. Freezing Base Model Layers:

Layers are freeze of the pre-trained ResNet50 model to prevent them from being updated during training. This is done by setting layer trainable = False for all layers in the base model.

5.1.5. Model Compilation:

The model is compiled using the Adam optimizer with a learning rate of 0.001. The loss function is set to 'categorical Cross entropy' , which is appropriate for multi-class classification. The model is configured to track accuracy as a metric during training.

5.1.6. Model Training:

The model is trained using the fit method with the training and validation data generated by train batches and valid batches.[9] Training continues for 30 epochs with a verbosity level of 2 (to display progress updates). It involves data preprocessing, model architecture design, freezing layers, model compilation, and training.

5.2 Price Prediction Model:

5.2.1. Data Preprocessing:

Perform data preprocessing by splitting and converting textual values in the 'Mileage,' 'Engine,' and 'Power' columns into numeric values. The 'Mileage,' 'Engine,' and 'Power' columns are cleaned and converted to numeric data types. The dataset is divided into features (X) and the target variable (y) for regression analysis.

5.2.2. Data Splitting:

The dataset is split into training and testing sets using the train test split function. This separation is essential to evaluate the model's performance.

5.2.3. Prediction Techniques:

A Linear Regression model is trained on the training data to predict the price of items. Predictions are made on the test set, and the model's performance is assessed using metrics like Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and R-squared score (R2) [10]. Lasso and Ridge regression models are fitted to the training data, and predictions are made on the test set. Similar to Linear Regression, performance metrics are calculated for both Lasso and Ridge regression models[11]. A Random Forest Regressor is trained with 100 estimators on the training data to predict item prices. Predictions are evaluated using the same metrics as other models. A Decision Tree Regressor is trained on the training data and used to predict item prices on the test set. The model's performance is assessed using the specified metrics[12]. An XGBoost Regressor is trained with default hyperparameters. A Grid Search Cross-Validation is performed to search for the best combination of hyperparameters that minimizes the mean squared error. The best hyperparameters and the corresponding best score are printed.

5.2.4. Final Model Selection and Hyperparameter Tuning:

Hyperparameters are crucial for configuring and fine-tuning the behavior of different regression models. Hyperparameter tuning is the process of finding the optimal values for these parameters, which can significantly impact the model's performance and generalization capabilities. Grid search and cross-validation are used in the code to identify the best hyperparameters for the corresponding models. An XGBoost Regressor model is initialized with the best hyperparameters obtained from the grid search. This model is then trained on the entire training dataset.

5. CONCLUSION

In this study, deep learning methods are used in a new way to create a system that can find damage to cars and guess how much it will cost to fix. The method used uses the VGG-16 and ResNet50 architectures to pull out important features from pictures of cars, which makes training a prediction model easier. The technology was able to accurately guess prices 90% of the time and find damage to cars 95% of the time. The suggested method could be used by both people and insurance companies to make the process of reviewing cases and getting a rough idea of how much it will cost to fix a damaged car go more quickly. The potential for enhancement lies in the utilization of a more extensive and heterogeneous dataset for the purpose of training, as well as the integration of supplementary variables into the model for price prediction.

6. REFERENCES

[1] B. Hamad, M. M. Kamel, and D. K. Du, "Vehicle detection in unconstrained aerial videos using deep learning," in Proceedings of the European Conference on Computer Vision (ECCV), 2018.

[2] R. Zhang, X. Yao, Z. He, and X. Li, "A survey of vehicle detection in unmanned aerial vehicles," IEEE Access, 2020.

[3] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? The KITTI vision benchmark suite," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2012.

[4] S. Zhang, X. Song, Y. Cai, and X. Zhang, "Vehicle detection in aerial images using deep learning and improved YOLOv3," Remote Sensing, 2020.

[5] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu, "Human3.6M: Large scale datasets and predictive methods for 3D human sensing in natural environments," IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2014.

[6] James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). An Introduction to Statistical Learning: with Applications in R. Springer.

[7] Hastie, T., Tibshirani, R., & Friedman, J. (2009). The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Springer.

[8] Brownlee, J. (2018). Introduction to Time Series Forecasting with Python. Machine Learning Mastery.

[9] Hyndman, R. J., & Athanasopoulos, G. (2018). Forecasting: principles and practice. OTexts.

[10] Chen, J., Song, L., Chen, L., & Lin, Q. (2020). Airbnb price prediction: Machine learning and econometric models. Tourism Management, 81, 104157.

[11] Osanaiye, O., Arshad, J., & Akhunzada, A. (2020). A comprehensive review of machine learning models for price prediction. IEEE Access, 8, 165227-16526

[12] Jeffrey de Deijn. 2018. Automatic Car Damage Recognition using Convolutional Neural Networks. (2018).

[13] Srikanth Tammina. [n.d.]. Transfer learning using VGG-16 with Deep Convolutional Neural Network for ClassifyingImages. ([n. d.]).

[14] Ranjodh Singh, Meghna P Ayyar, Tata Sri Pavan, Sandeep Gosain, and Rajiv Ratn Shah. 2019. Automating Car Insurance Claims Using Deep Learning Techniques. In 2019 IEEE Fifth International Conference on Multimedia Big Data (BigMM). IEEE, 199--207.

[15] Kalpesh Patil, Mandar Kulkarni, Anand Sriraman, and Shirish Karande. 2017. Deep learning based car damage classification. In 2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA). IEEE, 50--54.